



Australian Government

Patent Office
Canberra

I, JANENE PEISKER, MANAGER EXAMINATION SUPPORT AND SALES
hereby certify that annexed is a true copy of the Provisional specification in
connection with Application No. 2003901617 for a patent by SILVERBROOK
RESEARCH PTY LTD. as filed on 07 April 2003.

WITNESS my hand this
Sixth day of January 2006

A handwritten signature in black ink, appearing to read "J. Peisker".

JANENE PEISKER
MANAGER EXAMINATION SUPPORT
AND SALES

BEST AVAILABLE COPY



CERTIFIED COPY OF
PRIORITY DOCUMENT

METHODS AND SYSTEMS FOR OBJECT IDENTIFICATION AND INTERACTION

FIELD OF INVENTION

This invention relates to unique object identification and, in particular, to methods and systems for identifying and interacting with objects.

CO-PENDING APPLICATIONS

Various methods, systems and apparatus relating to the present invention are disclosed in the following US applications co-filed by the applicant or assignee of the present invention:

USSN ____/______ entitled "Orientation-Indicating Cyclic Position Codes" (docket number NPT019US);

USSN ____/______ entitled "Symmetric Tags" (docket number NPT023US);

USSN ____/______ entitled "Methods and Systems for Object Identification and Interaction" (docket number NIR001US);

USSN ____/______ entitled "Methods and Systems for Object Identification and Interaction" (docket number NIR003US); and

USSN ____/______ entitled "Methods and Systems for Object Identification and Interaction" (docket number NIR004US).

The disclosures of these co-filed applications are incorporated herein by cross-reference.

Various methods, systems and apparatus relating to the present invention are disclosed in the following application filed by the applicant or assignee of the present invention on 7 April 2003: Australian Provisional Application No _____ entitled "Methods and Systems for Object Identification and Interaction" (docket number NIR002AU). The disclosures of this co-pending application are incorporated herein by cross-reference.

Various methods, systems and apparatus relating to the present invention are disclosed in the following application filed by the applicant or assignee of the present invention on 17 February 2003: Australian Provisional Application No _____ entitled "Methods, Systems and Apparatus (NPS041)" (docket number NPS041). The disclosures of this co-pending application are incorporated herein by cross-reference.

Various methods, systems and apparatus relating to the present invention are disclosed in the following co-pending PCT applications filed by the applicant or assignee of

the present invention on 4 December 2002: USSN ____/______ entitled "Rotationally Symmetric Tags" (docket number NPT020US). The disclosures of these co-pending applications are incorporated herein by cross-reference.

Various methods, systems and apparatus relating to the present invention are disclosed in the following co-pending PCT applications filed by the applicant or assignee of the present invention on 22 November 2002:

PCT/AU02/01572 and PCT/AU/02/01573.

The disclosures of these co-pending applications are incorporated herein by cross-reference.

Various methods, systems and apparatus relating to the present invention are disclosed in the following co-pending PCT applications filed by the applicant or assignee of the present invention on 15 October 2002:

PCT/AU02/01391, PCT/AU02/01392, PCT/AU02/01393, PCT/AU02/01394 and PCT/AU02/01395.

The disclosures of these co-pending applications are incorporated herein by cross-reference.

Various methods, systems and apparatus relating to the present invention are disclosed in the following co-pending PCT applications filed by the applicant or assignee of the present invention on 26 November 2001:

PCT/AU01/01527, PCT/AU01/01528, PCT/AU01/01529, PCT/AU01/01530 and PCT/AU01/01531.

The disclosures of these co-pending applications are incorporated herein by cross-reference.

Various methods, systems and apparatus relating to the present invention are disclosed in the following co-pending PCT applications filed by the applicant or assignee of the present invention on 11 October 2001: PCT/AU01/01274. The disclosures of these co-pending applications are incorporated herein by cross-reference.

Various methods, systems and apparatus relating to the present invention are disclosed in the following co-pending PCT applications filed by the applicant or assignee of the present invention on 14 August 2001: PCT/AU01/00996. The disclosures of these co-pending applications are incorporated herein by cross-reference.

Various methods, systems and apparatus relating to the present invention are disclosed in the following co-pending US applications filed by the applicant or assignee of the present invention on 27 November 2000:

09/721895, 09/721894, 09/722174, 09/721896, 09/722148, 09/722146, 09/721861, 09/721892, 09/722171, 09/721858, 09/722142, 09/722087, 09/722141, 09/722175, 09/722147, 09/722172, 09/721893, 09/722088, 09/721862, 09/721856, 09/721857, 09/721859, 09/721860

The disclosures of these co-pending applications are incorporated herein by cross-reference.

Various methods, systems and apparatus relating to the present invention are disclosed in the following co-pending US applications filed by the applicant or assignee of the present invention on 20 October 2000:

09/693415, 09/693219, 09/693280, 09/693515, 09/693705, 09/693647, 09/693690, 09/693593, 09/693216, 09/693341, 09/696473, 09/696514, 09/693301, 09/693388, 09/693704 09/693510, 09/693336, 09/693335

The disclosures of these co-pending US applications are incorporated herein by cross-reference.

Various methods, systems and apparatus relating to the present invention are disclosed in the following co-pending US applications filed by the applicant or assignee of the present invention on 15 September 2000:

09/663579, 09/669599, 09/663701, 09/663640

The disclosures of these co-pending applications are incorporated herein by cross-reference.

Various methods, systems and apparatus relating to the present invention are disclosed in the following co-pending US applications filed by the applicant or assignee of the present invention on 30 June 2000:

09/609139, 09/608970, 09/609039, 09/607852, 09/607656, 09/609132, 09/609303, 09/610095, 09/609596, 09/607843, 09/607605, 09/608178, 09/609553, 09/609233, 09/609149, 09/608022, 09/609232, 09/607844, 09/607657, 09/608920, 09/607985, 09/607990 09/607196, 09/606999

The disclosures of these co-pending applications are incorporated herein by cross-reference.

Various methods, systems and apparatus relating to the present invention are disclosed in the following co-pending US applications filed by the applicant or assignee of the present invention on 23 May 2000:

09/575197, 09/575195, 09/575159, 09/575132, 09/575123, 09/575148,
09/575130, 09/575165, 09/575153, 09/575118, 09/575131, 09/575116,
09/575144, 09/575139, 09/575186, 09/575185, 09/575191, 09/575145;
09/575192, 09/575181, 09/575193, 09/575156, 09/575183, 09/575160,
09/575150, 09/575169, 09/575184, 09/575128, 09/575180, 09/575149,
09/575179, 09/575187, 09/575155, 09/575133, 09/575143, 09/575196,
09/575198, 09/57578 09/575164, 09/575146, 09/575174, 09/575163,
09/575168, 09/575154, 09/575129, 09/575124, 09/575188, 09/575189,
09/575162, 09/575172, 09/575170, 09/575171, 09/575161, 09/575141,
09/575125, 09/575142, 09/575140, 09/575190, 09/575138, 09/575126,
09/575127, 09/575158, 09/575117, 09/575147, 09/575152, 09/575176,
09/575115, 09/575114, 09/575113, 09/575112, 09/575111, 09/575108,
09/575109, 09/575110

The disclosures of these co-pending applications are incorporated herein by cross-reference.

BACKGROUND

For the purposes of automatic identification, a product item is commonly identified by a 12-digit Universal Product Code (UPC), encoded machine-readably in the form of a printed bar code.

Within supply chain management, there is considerable interest in expanding or replacing the UPC scheme to allow individual product items to be uniquely identified and thereby tracked. Individual item tagging can reduce "shrinkage" due to lost, stolen or spoiled goods, improve the efficiency of demand-driven manufacturing and supply, facilitate the profiling of product usage, and improve the customer experience.

There are two main contenders for individual item tagging: optical tags in the form of two-dimensional bar codes, and radio frequency identification (RFID) tags. Optical tags have the advantage of being inexpensive, but require optical line-of-sight for reading. RFID tags have the advantage of supporting omnidirectional reading, but are comparatively expensive.

SUMMARY OF THE INVENTION

According to a first aspect of the present invention there is disclosed a method of facilitating an interaction between a user and a product item, the product item having an identity and the method including the steps of:

providing the user with an interface surface associated with the product item and containing information relating to the product item, the interface surface having disposed thereon coded data indicative of the identity of the product item and of a plurality of reference points of the interface surface;

receiving, in a computer system, indicating data from a sensing device regarding the identity of the product item and a position of the sensing device relative to the interface surface, the sensing device, when placed in an operative position relative to the interface surface, sensing at least some of the coded data in the vicinity of the sensing device and generating the indicating data using at least some of the sensed coded data; and

facilitating, in the computer system and with reference to the indicating data, the interaction between the user and the product item.

Preferably, the interaction is associated with at least one zone of the interface surface and in which the method includes identifying, in the computer system and from the zone relative to which the sensing device is located, the interaction.

Preferably, the method includes:

receiving, in the computer system, data regarding movement of the sensing device relative to the interface surface, the sensing device sensing its movement relative to the interface surface using at least some of the coded data; and

identifying, in the computer system and from the movement being at least partially within the at least one zone, the interaction.

According to a second aspect of the present invention there is disclosed a method of facilitating an interaction between a user and a product item, the product item having an identity and the method including the steps of:

providing the user with an interface surface associated with the product item, the interface surface containing information relating to the product item and having disposed thereon coded data indicative of the identity of the product item and a parameter of the interaction;

receiving, in a computer system and from a sensing device, indicating data regarding the identity of both the product item and the parameter of the interaction, and movement data regarding movement of the sensing device relative to the interface surface, the sensing device, when moved relative to the interface surface, sensing at least some of the coded data and generating the indicating data and the movement data using at least some of the sensed coded data; and

interpreting, in the computer system, the indicating data and the movement data as it relates to the interaction.

According to a third aspect of the present invention there is disclosed a method of facilitating an interaction between a user and a product item, the product item having an identity and the method including the steps of:

providing the user with an interface surface associated with the product item, the interface surface containing information relating to the product item and having disposed thereon coded data indicative of the identity of the product item;

receiving, in a computer system, user data from a sensing device regarding an identity of the user and indicating data regarding the identity of the product item, the sensing device containing the user data and, when placed in an operative position relative to the interface surface, sensing at least some of the coded data in the vicinity of the sensing device and generating the indicating data using at least some of the sensed coded data; and

facilitating, in the computer system and with reference to the user data and the indicating data, the interaction between the user and the product item.

Preferably, the coded data is also indicative of a parameter of the interaction, and the method includes receiving, in the computer system, indicating data from the sensing device regarding the parameter of the interaction, the sensing device generating the indicating data using at least some of the sensed coded data.

Preferably, the method includes receiving, in the computer system, data from the sensing device regarding movement of the sensing device relative to the interface surface, the sensing device generating data regarding its own movement relative to the interface surface.

Preferably, the interaction is selected from the group comprising:

- (a) providing product information about the product item to the user;
- (b) recording a purchase transaction relating to the product item;
- (c) recording a potential purchase transaction relating to the product item;
- (d) providing comparison information to the user, the comparison information comparing the product information about the product item with product information about another product item;
- (e) playing a game associated with the product item; and
- (f) conducting a competition in relation to the product item.

Preferably, the product information comprises information relating to any one of the product item's:

- (a) cost;
- (b) contents;
- (c) weight;
- (d) place of origin;
- (e) manufacturer;
- (f) date of manufacture;
- (g) date of packaging;
- (h) use-by date;
- (i) current owner; and
- (j) dimensions.

Preferably, the interface surface is selected from the group comprising:

- (a) a label;
- (b) a package; and
- (c) a surface of the product item itself.

Preferably, the coded data is substantially invisible to the average unaided human.

Preferably, the coded data is disposed over a substantial portion of the interface surface. More preferably, the coded data is disposed over more than 20% of the interface surface. Even more preferably, the coded data is disposed over more than 90% of the interface surface.

Preferably, the method further comprises identifying, in the computer system, that the user has entered handwritten text data by means of the sensing device and effecting, in the computer system, an operation associated with the handwritten text data.

Preferably, the method includes converting, in the computer system, the handwritten text data to computer text.

Preferably, the method further comprises identifying, in the computer system, that the user has entered a handwritten signature by means of the sensing device and effecting, in the computer system, an operation associated with the handwritten signature.

Preferably, the method includes verifying, in the computer system, that the signature is that of the user.

Preferably, the operation associated with the handwritten signature is associated with payment authorization.

Preferably, the method further comprises identifying, in the computer system, that the user has entered a hand-drawn picture by means of the sensing device and effecting, in the computer system, an operation associated with the hand-drawn picture.

Preferably, a portion of the coded data is superimposed with a visual graphic, the visual graphic being at least partially indicative, to the user, of the interaction.

Preferably, the sensing device contains an identifier which imparts a unique identity to the sensing device and identifies it as belonging to the user and in which the method includes monitoring, in the computer system, the identifier.

Preferably, the method includes providing all required information relating to the interaction in the interface surface to eliminate the need for a separate display device.

In one form, the coded data is disposed on the interface surface in accordance with a layout, the layout having at least order n rotational symmetry, where n is at least two, the

layout encoding an orientation codeword comprising a sequence of an integer multiple m of n symbols, where m is one or more, each encoded symbol being distributed at n locations about a center of rotational symmetry of the layout such that decoding the symbols at each of the n orientations of the layout produces n representations of the orientation codeword, each representation comprising a different cyclic shift of the orientation codeword and being indicative of the degree of rotation of the layout, and wherein the orientation codeword is fault tolerant.

In another form, the coded data is disposed on the interface surface in accordance with a layout, the layout having at least order n rotational symmetry, where n is at least two, the layout including n identical sub-layouts rotated $1/n$ revolutions apart about a center of rotational symmetry of the layout, the coded data disposed in accordance with each sub-layout including rotation-indicating data that distinguishes the rotation of that sub-layout from the rotation of at least one other sub-layout within the layout.

In a further form, the coded data is disposed on the interface surface in accordance with a layout having n -fold rotational symmetry, where n is at least two, the layout including n identical first sub-layouts rotated $1/n$ revolutions apart about a center of rotational symmetry of the layout, the coded data disposed in accordance with each first sub-layout including rotation-indicating data that distinguishes the rotation of that first sub-layout from the rotation of at least one other first sub-layout within the layout.

According to a fourth aspect of the present invention there is disclosed a system for facilitating an interaction between a user and a product item, the product item having an identity and the system including:

a interface surface associated with the product item and containing information relating to the product item, the interface surface having disposed thereon coded data indicative of the identity of the product item and of a plurality of reference points of the interface surface; and

a computer system adapted to facilitate the interaction in response to receiving indicating data from a sensing device, the indicating data being indicative of the identity of the product item and of a position of the sensing device relative to the interface surface, the sensing device, when placed in an operative position relative to the interface surface, sensing at least some of the coded data in the vicinity of the sensing device and generating the indicating data using at least some of the sensed coded data.

Preferably, the interaction is associated with at least one zone of the interface surface.

Preferably, the system includes the sensing device, the sensing device sensing its movement relative to the interface surface using at least some of the coded data.

According to a fifth aspect of the present invention there is disclosed a system for facilitating an interaction between a user and a product item, the product item having an identity and the system including

an interface surface associated with the product item, the interface surface containing information relating to the product item and having disposed thereon coded data indicative of the identity of the product item and of a parameter of the interaction; and

a computer system for receiving, from a sensing device, indicating data regarding the identity of both the product item and the parameter of the interaction, and movement data regarding movement of the sensing device relative to the interface surface, and for interpreting the movement of the sensing device as it relates to the interaction, the sensing device, when moved relative to the interface surface, sensing at least some of the coded data and generating the indicating data and the movement data using at least some of the sensed coded data.

According to a sixth aspect of the present invention there is disclosed a system for facilitating an interaction between a user and a product item, the product item having an identity and the system including:

a interface surface associated with the product item, the interface surface containing information relating to the product item and having disposed thereon coded data indicative of the identity of the product item; and

a computer system adapted to receive from a sensing device user data regarding an identity of the user and indicating data regarding the identity of the product item, and for facilitating, with reference to the user data and the indicating data, the interaction between the user and the product item, the sensing device containing the user data and, when placed in an operative position relative to the interface surface, sensing at least some of the coded data in the vicinity of the sensing device and generating the indicating data using at least some of the sensed coded data.

Preferably, the coded data is also indicative of a parameter of the interaction, the computer system receiving indicating data from the sensing device regarding the parameter of the interaction, and the sensing device generating the indicating data using at least some of the coded data.

Preferably, the system includes the sensing device, the sensing device sensing its movement relative to the interface surface.

Preferably, the interaction is selected from the group comprising:

- (a) providing product information about the product item to the user;
- (b) recording a purchase transaction relating to the product item;
- (c) recording a potential purchase transaction relating to the product item;
- (d) providing comparison information to the user, the comparison information comparing the product information about the product item with product information about another product item;
- (e) playing a game associated with the product item; and
- (f) conducting a competition in relation to the product item.

Preferably, the product information comprises information relating to any one of the product item's:

- (a) cost;
- (b) contents;
- (c) weight;
- (d) place of origin;
- (e) manufacturer;
- (f) date of manufacture;
- (g) date of packaging;
- (h) use-by date;
- (i) current owner; and
- (j) dimensions.

Preferably, the interface surface is selected from the group comprising:

- (a) a label;
- (b) a package; and
- (c) a surface of the product item itself.

Preferably, the coded data is substantially invisible to the average unaided human.

Preferably, the coded data is disposed over a substantial portion of the interface surface. More preferably, the coded data is disposed over more than 20% of the interface

surface. Even more preferably, the coded data is disposed over more than 90% of the interface surface.

Preferably, the sensing device includes a marking nib.

Preferably, the sensing device contains an identifier which imparts a unique identity to the sensing device and identifies it as belonging to the user.

In one form, the coded data is disposed on the interface surface in accordance with a layout, the layout having at least order n rotational symmetry, where n is at least two, the layout encoding an orientation codeword comprising a sequence of an integer multiple m of n symbols, where m is one or more, each encoded symbol being distributed at n locations about a center of rotational symmetry of the layout such that decoding the symbols at each of the n orientations of the layout produces n representations of the orientation codeword, each representation comprising a different cyclic shift of the orientation codeword and being indicative of the degree of rotation of the layout, and wherein the orientation codeword is fault tolerant.

In another form, the coded data is disposed on the interface surface in accordance with a layout, the layout having at least order n rotational symmetry, where n is at least two, the layout including n identical sub-layouts rotated $1/n$ revolutions apart about a center of rotational symmetry of the layout, the coded data disposed in accordance with each sub-layout including rotation-indicating data that distinguishes the rotation of that sub-layout from the rotation of at least one other sub-layout within the layout.

In a further form, the coded data is disposed on the interface surface in accordance with a layout having n -fold rotational symmetry, where n is at least two, the layout including n identical first sub-layouts rotated $1/n$ revolutions apart about a center of rotational symmetry of the layout, the coded data disposed in accordance with each first sub-layout including rotation-indicating data that distinguishes the rotation of that first sub-layout from the rotation of at least one other first sub-layout within the layout.

According to a seventh aspect of the present invention there is disclosed an interactive product item adapted for interaction with a user via a sensing device and a computer system, and comprising:

a product item having an identity;

an interface surface associated with the product item and having disposed thereon information relating to the product item and coded data indicative of the identity of the product item.

Preferably, the coded data is further indicative of a plurality of reference points of the interface surface.

Preferably, the coded data is further indicative of a parameter of an interaction.

Preferably, the interface surface is selected from the group comprising:

- (a) a label;
- (b) a package; and
- (c) a surface of the product item itself.

Preferably, the coded data is disposed on the interface surface in accordance with a layout, the layout having at least order n rotational symmetry, where n is at least two, the layout encoding an orientation codeword comprising a sequence of an integer multiple m of n symbols, where m is one or more, each encoded symbol being distributed at n locations about a center of rotational symmetry of the layout such that decoding the symbols at each of the n orientations of the layout produces n representations of the orientation codeword, each representation comprising a different cyclic shift of the orientation codeword and being indicative of the degree of rotation of the layout, and wherein the orientation codeword is fault tolerant.

Alternatively, the coded data is disposed on the interface surface in accordance with a layout, the layout having at least order n rotational symmetry, where n is at least two, the layout including n identical sub-layouts rotated $1/n$ revolutions apart about a center of rotational symmetry of the layout, the coded data disposed in accordance with each sub-layout including rotation-indicating data that distinguishes the rotation of that sub-layout from the rotation of at least one other sub-layout within the layout.

In another form, the coded data is disposed on the interface surface in accordance with a layout having n -fold rotational symmetry, where n is at least two, the layout including n identical first sub-layouts rotated $1/n$ revolutions apart about a center of rotational symmetry of the layout, the coded data disposed in accordance with each first sub-layout including rotation-indicating data that distinguishes the rotation of that first sub-layout from the rotation of at least one other first sub-layout within the layout.

BRIEF DESCRIPTION OF THE DRAWINGS

Preferred and other embodiments of the invention will now be described, by way of non-limiting example only, with reference to the accompanying drawings, in which:

Figure 1 is a schematic of a the relationship between a sample printed netpage and its online page description;

Figure 2 is a schematic view of a interaction between a netpage pen, a Web terminal, a netpage printer, a netpage relay, a netpage page server, and a netpage application server, and a Web server;

Figure 3 illustrates a collection of netpage servers, Web terminals, printers and relays interconnected via a network;

Figure 4 is a schematic view of a high-level structure of a printed netpage and its online page description;

Figure 5a is a plan view showing the interleaving and rotation of the symbols of four codewords of the tag;

Figure 5b is a plan view showing a macrodot layout for the tag shown in Figure 5a;

Figure 5c is a plan view showing an arrangement of nine of the tags shown in Figures 5a and 5b, in which targets are shared between adjacent tags;

Figure 6 is a plan view showing a relationship between a set of the tags shown in Figure 6a and a field of view of a netpage sensing device in the form of a netpage pen;

Figure 7 is a flowchart of a tag image processing and decoding algorithm;

Figure 8 is a perspective view of a netpage pen and its associated tag-sensing field-of-view cone;

Figure 9 is a perspective exploded view of the netpage pen shown in Figure 8;

Figure 10 is a schematic block diagram of a pen controller for the netpage pen shown in Figures 8 and 9;

Figure 11 is a perspective view of a wall-mounted netpage printer;

Figure 12 is a section through the length of the netpage printer of Figure 11;

Figure 12a is an enlarged portion of Figure 12 showing a section of the duplexed print engines and glue wheel assembly;

Figure 13 is a detailed view of the ink cartridge, ink, air and glue paths, and print engines of the netpage printer of Figures 11 and 12;

Figure 14 is a schematic block diagram of a printer controller for the netpage printer shown in Figures 11 and 12;

Figure 15 is a schematic block diagram of duplexed print engine controllers and Memjet™ printheads associated with the printer controller shown in Figure 14;

Figure 16 is a schematic block diagram of the print engine controller shown in Figures 14 and 15;

Figure 17 is a perspective view of a single Memjet™ printing element, as used in, for example, the netpage printer of Figures 10 to 12;

Figure 18 is a schematic view of the structure of an item ID;

Figure 19 is a schematic view of the structure of an omnitag;

Figure 20 is a schematic view of a product item and object ownership and packaging hierarchy class diagram;

Figure 21 is a schematic view of a user class diagram;

Figure 22 is a schematic view of a printer class diagram;

Figure 23 is a schematic view of a pen class diagram;

Figure 24 is a schematic view of an application class diagram;

Figure 25 is a schematic view of a document and page description class diagram;

Figure 26 is a schematic view of a document and page ownership class diagram;

Figure 27 is a schematic view of a terminal element specialization class diagram;

Figure 28 is a schematic view of a static element specialization class diagram;

Figure 29 is a schematic view of a hyperlink element class diagram;

Figure 30 is a schematic view of a hyperlink element specialization class diagram;

Figure 31 is a schematic view of a hyperlinked group class diagram;

Figure 32 is a schematic view of a form class diagram;

Figure 33 is a schematic view of a digital ink class diagram;

Figure 34 is a schematic view of a field element specialization class diagram;

Figure 35 is a schematic view of a checkbox field class diagram;

Figure 36 is a schematic view of a text field class diagram;

Figure 37 is a schematic view of a signature field class diagram;

Figure 38 is a flowchart of an input processing algorithm;

Figure 38a is a detailed flowchart of one step of the flowchart of Figure 38;

Figure 39 is a schematic view of a page server command element class diagram;

Figure 40 is a schematic view of a subscription delivery protocol;

Figure 41 is a schematic view of a hyperlink request class diagram;

Figure 42 is a schematic view of a hyperlink activation protocol;

Figure 43 is a schematic view of a form submission protocol;

Figure 44 shows a triangular macrodot packing with a four-bit symbol unit outlined, for use with an embodiment of the invention;

Figure 45 shows a square macrodot packing with a four-bit symbol unit outlined, for use with an embodiment of the invention such as that described in relation to Figures 5a to 5c;

Figure 46 shows a one-sixth segment of an hexagonal tag, with the segment containing a maximum of 11 four-bit symbols with the triangular macrodot packing shown in Figure 44;

Figure 47 shows a one-quarter segment of a square tag, with the segment containing a maximum of 15 four-bit symbols with the square macrodot packing shown in Figure 45;

Figure 48 shows a logical layout of a hexagonal tag using the tag segment of Figure 47, with six interleaved 2^4 -ary $(11, k)$ codewords;

Figure 49 shows the macrodot layout of the hexagonal tag of Figure 48;

Figure 50 shows an arrangement of seven abutting tags of the design of Figures 48 and 49, with shared targets;

Figure 51 shows a logical layout of an alternative hexagonal tag using the tag segment of Figure 47, with three interleaved 2^4 -ary $(9, k)$ codewords and three interleaved three-symbol fragments of three distributed 2^4 -ary $(9, k)$ codewords;

Figure 52 shows the logical layout of an orientation-indicating cyclic position codeword of the hexagonal tag of Figure 51;

Figure 53 shows three adjacent tags of type P, Q and R, each with the layout of the tag of Figure 51, containing a complete set of distributed codewords;

Figure 54 shows the logical layout of yet another alternative hexagonal tag using the tag segment of Figure 47, with one local 2^4 -ary (12, k) codeword, interleaved with eighteen 3-symbol fragments of eighteen distributed 2^4 -ary (9, k) codewords;

Figure 55 shows the logical layout of the hexagonal tag of Figure 54, re-arranged to show the distributed 3-symbol fragments which contribute to the same codewords;

Figure 56 is a schematic view of a physical product item and its online description; and

Figure 57 is a schematic view of the interaction between a product item, a fixed product scanner, a hand-held product scanner, a scanner relay, a product server, and a product application server.

DESCRIPTION OF PREFERRED AND OTHER EMBODIMENTS

Note: Memjet™ is a trade mark of Silverbrook Research Pty Ltd, Australia.

In the preferred embodiment, the invention is configured to work with the netpage networked computer system, a detailed overview of which follows. It will be appreciated that not every implementation will necessarily embody all or even most of the specific details and extensions discussed below in relation to the basic system. However, the system is described in its most complete form to reduce the need for external reference when attempting to understand the context in which the preferred embodiments and aspects of the present invention operate.

In brief summary, the preferred form of the netpage system employs a computer interface in the form of a mapped surface, that is, a physical surface which contains references to a map of the surface maintained in a computer system. The map references can be queried by an appropriate sensing device. Depending upon the specific implementation, the map references may be encoded visibly or invisibly, and defined in such a way that a local query on the mapped surface yields an unambiguous map reference both within the map and among different maps. The computer system can contain information about features on the mapped surface, and such information can be retrieved based on map references supplied by a sensing device used with the mapped surface. The information thus retrieved can take the form of actions which are initiated by the computer system on behalf of the operator in response to the operator's interaction with the surface features.

In its preferred form, the netpage system relies on the production of, and human interaction with, netpages. These are pages of text, graphics and images printed on ordinary paper, but which work like interactive web pages. Information is encoded on each page using ink which is substantially invisible to the unaided human eye. The ink, however, and thereby the coded data, can be sensed by an optically imaging pen and transmitted to the netpage system.

In the preferred form, active buttons and hyperlinks on each page can be clicked with the pen to request information from the network or to signal preferences to a network server. In one embodiment, text written by hand on a netpage is automatically recognized and converted to computer text in the netpage system, allowing forms to be filled in. In other embodiments, signatures recorded on a netpage are automatically verified, allowing e-commerce transactions to be securely authorized.

As illustrated in Figure 1, a printed netpage 1 can represent an interactive form which can be filled in by the user both physically, on the printed page, and "electronically", via communication between the pen and the netpage system. The example shows a

"Request" form containing name and address fields and a submit button. The netpage consists of graphic data 2 printed using visible ink, and coded data 3 printed as a collection of tags 4 using invisible ink. The corresponding page description 5, stored on the netpage network, describes the individual elements of the netpage. In particular it describes the type and spatial extent (zone) of each interactive element (i.e. text field or button in the example), to allow the netpage system to correctly interpret input via the netpage. The submit button 6, for example, has a zone 7 which corresponds to the spatial extent of the corresponding graphic 8.

As illustrated in Figure 2, the netpage pen 101, a preferred form of which is shown in Figures 8 and 9 and described in more detail below, works in conjunction with a personal computer (PC), Web terminal 75, or a netpage printer 601. The netpage printer is an Internet-connected printing appliance for home, office or mobile use. The pen is wireless and communicates securely with the netpage network via a short-range radio link 9. Short-range communication is relayed to the netpage network by a local relay function which is either embedded in the PC, Web terminal or netpage printer, or is provided by a separate relay device 44. The relay function can also be provided by a mobile phone or other device which incorporates both short-range and longer-range communications functions.

In an alternative embodiment, the netpage pen utilises a wired connection, such as a USB or other serial connection, to the PC, Web terminal, netpage printer or relay device.

The netpage printer 601, a preferred form of which is shown in Figures 11 to 13 and described in more detail below, is able to deliver, periodically or on demand, personalized newspapers, magazines, catalogs, brochures and other publications, all printed at high quality as interactive netpages. Unlike a personal computer, the netpage printer is an appliance which can be, for example, wall-mounted adjacent to an area where the morning news is first consumed, such as in a user's kitchen, near a breakfast table, or near the household's point of departure for the day. It also comes in tabletop, desktop, portable and miniature versions.

Netpages printed at their point of consumption combine the ease-of-use of paper with the timeliness and interactivity of an interactive medium.

As shown in Figure 2, the netpage pen 101 interacts with the coded data on a printed netpage 1 (or product item 201) and communicates the interaction via a short-range radio link 9 to a relay. The relay sends the interaction to the relevant netpage page server 10 for interpretation. In appropriate circumstances, the page server sends a corresponding message to application computer software running on a netpage application server 13. The application server may in turn send a response which is printed on the originating printer.

In an alternative embodiment, the PC, Web terminal, netpage printer or relay device may communicate directly with local or remote application software, including a local or remote Web server. Relatedly, output is not limited to being printed by the netpage printer. It can also be displayed on the PC or Web terminal, and further interaction can be screen-based rather than paper-based, or a mixture of the two.

The netpage system is made considerably more convenient in the preferred embodiment by being used in conjunction with high-speed microelectromechanical system (MEMS) based inkjet (Memjet™) printers. In the preferred form of this technology, relatively high-speed and high-quality printing is made more affordable to consumers. In its preferred form, a netpage publication has the physical characteristics of a traditional news-magazine, such as a set of letter-size glossy pages printed in full color on both sides, bound together for easy navigation and comfortable handling.

The netpage printer exploits the growing availability of broadband Internet access. Cable service is available to 95% of households in the United States, and cable modem service offering broadband Internet access is already available to 20% of these. The netpage printer can also operate with slower connections, but with longer delivery times and lower image quality. Indeed, the netpage system can be enabled using existing consumer inkjet and laser printers, although the system will operate more slowly and will therefore be less acceptable from a consumer's point of view. In other embodiments, the netpage system is hosted on a private intranet. In still other embodiments, the netpage system is hosted on a single computer or computer-enabled device, such as a printer.

Netpage publication servers 14 on the netpage network are configured to deliver print-quality publications to netpage printers. Periodical publications are delivered automatically to subscribing netpage printers via pointcasting and multicasting Internet protocols. Personalized publications are filtered and formatted according to individual user profiles.

A netpage printer can be configured to support any number of pens, and a pen can work with any number of netpage printers. In the preferred implementation, each netpage pen has a unique identifier. A household may have a collection of colored netpage pens, one assigned to each member of the family. This allows each user to maintain a distinct profile with respect to a netpage publication server or application server.

A netpage pen can also be registered with a netpage registration server 11 and linked to one or more payment card accounts. This allows e-commerce payments to be securely authorized using the netpage pen. The netpage registration server compares the signature captured by the netpage pen with a previously registered signature, allowing it to authenticate the user's identity to an e-commerce server. Other biometrics can also be used

to verify identity. A version of the netpage pen includes fingerprint scanning, verified in a similar way by the netpage registration server.

Although a netpage printer may deliver periodicals such as the morning newspaper without user intervention, it can be configured never to deliver unsolicited junk mail. In its preferred form, it only delivers periodicals from subscribed or otherwise authorized sources. In this respect, the netpage printer is unlike a fax machine or e-mail account which is visible to any junk mailer who knows the telephone number or email address.

1 NETPAGE SYSTEM ARCHITECTURE

Each object model in the system is described using a Unified Modeling Language (UML) class diagram. A class diagram consists of a set of object classes connected by relationships, and two kinds of relationships are of interest here: associations and generalizations. An association represents some kind of relationship between objects, i.e. between instances of classes. A generalization relates actual classes, and can be understood in the following way: if a class is thought of as the set of all objects of that class, and class A is a generalization of class B, then B is simply a subset of A. The UML does not directly support second-order modelling - i.e. classes of classes.

Each class is drawn as a rectangle labelled with the name of the class. It contains a list of the attributes of the class, separated from the name by a horizontal line, and a list of the operations of the class, separated from the attribute list by a horizontal line. In the class diagrams which follow, however, operations are never modelled.

An association is drawn as a line joining two classes, optionally labelled at either end with the multiplicity of the association. The default multiplicity is one. An asterisk (*) indicates a multiplicity of "many", i.e. zero or more. Each association is optionally labelled with its name, and is also optionally labelled at either end with the role of the corresponding class. An open diamond indicates an aggregation association ("is-part-of"), and is drawn at the aggregator end of the association line.

A generalization relationship ("is-a") is drawn as a solid line joining two classes, with an arrow (in the form of an open triangle) at the generalization end.

When a class diagram is broken up into multiple diagrams, any class which is duplicated is shown with a dashed outline in all but the main diagram which defines it. It is shown with attributes only where it is defined.

1.1 NETPAGES

Netpages are the foundation on which a netpage network is built. They provide a paper-based user interface to published information and interactive services.

A netpage consists of a printed page (or other surface region) invisibly tagged with references to an online description of the page. The online page description is maintained persistently by a netpage page server. The page description describes the visible layout and content of the page, including text, graphics and images. It also describes the input elements on the page, including buttons, hyperlinks, and input fields. A netpage allows markings made with a netpage pen on its surface to be simultaneously captured and processed by the netpage system.

Multiple netpages can share the same page description. However, to allow input through otherwise identical pages to be distinguished, each netpage is assigned a unique page identifier. This page ID has sufficient precision to distinguish between a very large number of netpages.

Each reference to the page description is encoded in a printed tag. The tag identifies the unique page on which it appears, and thereby indirectly identifies the page description. The tag also identifies its own position on the page. Characteristics of the tags are described in more detail below.

Tags are printed in infrared-absorptive ink on any substrate which is infrared-reflective, such as ordinary paper. Near-infrared wavelengths are invisible to the human eye but are easily sensed by a solid-state image sensor with an appropriate filter.

A tag is sensed by an area image sensor in the netpage pen, and the tag data is transmitted to the netpage system via the nearest netpage printer. The pen is wireless and communicates with the netpage printer via a short-range radio link. Tags are sufficiently small and densely arranged that the pen can reliably image at least one tag even on a single click on the page. It is important that the pen recognize the page ID and position on every interaction with the page, since the interaction is stateless. Tags are error-correctably encoded to make them partially tolerant to surface damage.

The netpage page server maintains a unique page instance for each printed netpage, allowing it to maintain a distinct set of user-supplied values for input fields in the page description for each printed netpage.

The relationship between the page description, the page instance, and the printed netpage is shown in Figure 4. The printed netpage may be part of a printed netpage docu-

ment 45. The page instance is associated with both the netpage printer which printed it and, if known, the netpage user who requested it.

As shown in Figure 4, one or more netpages may also be associated with a physical object such as a product item, for example when printed onto the product item's label, packaging, or actual surface.

1.2 NETPAGE TAGS

1.2.1 Tag Data Content

In a preferred form, each tag identifies the region in which it appears, and the location of that tag within the region. A tag may also contain flags which relate to the region as a whole or to the tag. One or more flag bits may, for example, signal a tag sensing device to provide feedback indicative of a function associated with the immediate area of the tag, without the sensing device having to refer to a description of the region. A netpage pen may, for example, illuminate an "active area" LED when in the zone of a hyperlink.

As will be more clearly explained below, in a preferred embodiment, each tag contains an easily recognized invariant structure which aids initial detection, and which assists in minimizing the effect of any warp induced by the surface or by the sensing process. The tags preferably tile the entire page, and are sufficiently small and densely arranged that the pen can reliably image at least one tag even on a single click on the page. It is important that the pen recognize the page ID and position on every interaction with the page, since the interaction is stateless.

In a preferred embodiment, the region to which a tag refers coincides with an entire page, and the region ID encoded in the tag is therefore synonymous with the page ID of the page on which the tag appears. In other embodiments, the region to which a tag refers can be an arbitrary subregion of a page or other surface. For example, it can coincide with the zone of an interactive element, in which case the region ID can directly identify the interactive element.

In the preferred form, each tag contains 120 bits of information. The region ID is typically allocated up to 100 bits, the tag ID at least 16 bits, and the remaining bits are allocated to flags etc. Assuming a tag density of 64 per square inch, a 16-bit tag ID supports a region size of up to 1024 square inches. Larger regions can be mapped continuously without increasing the tag ID precision simply by using abutting regions and maps. The 100-bit region ID allows 2^{100} ($\sim 10^{30}$ or a million trillion trillion) different regions to be uniquely identified.

1.2.2 Tag Data Encoding

In one embodiment, the 120 bits of tag data are redundantly encoded using a (15, 5) Reed-Solomon code. This yields 360 encoded bits consisting of 6 codewords of 15 4-bit symbols each. The (15, 5) code allows up to 5 symbol errors to be corrected per codeword, i.e. it is tolerant of a symbol error rate of up to 33% per codeword.

Each 4-bit symbol is represented in a spatially coherent way in the tag, and the symbols of the six codewords are interleaved spatially within the tag. This ensures that a burst error (an error affecting multiple spatially adjacent bits) damages a minimum number of symbols overall and a minimum number of symbols in any one codeword, thus maximising the likelihood that the burst error can be fully corrected.

Any suitable error-correcting code can be used in place of a (15, 5) Reed-Solomon code, for example: a Reed-Solomon code with more or less redundancy, with the same or different symbol and codeword sizes; another block code; or a different kind of code, such as a convolutional code (see, for example, Stephen B. Wicker, Error Control Systems for Digital Communication and Storage, Prentice-Hall 1995, the contents of which are herein incorporated by reference thereto).

In order to support "single-click" interaction with a tagged region via a sensing device, the sensing device must be able to see at least one entire tag in its field of view no matter where in the region or at what orientation it is positioned. The required diameter of the field of view of the sensing device is therefore a function of the size and spacing of the tags.

1.2.3 Tag Structure

Figure 5a shows a tag 4, in the form of tag 726 with four perspective targets 17. The tag 726 represents sixty 4-bit Reed-Solomon symbols 747 (see description of Figures 44 to 46 below for discussion of symbols), for a total of 240 bits. The tag represents each "one" bit by the presence of a mark 748, referred to as a macrodot, and each "zero" bit by the absence of the corresponding macrodot. Figure 5c shows a square tiling 728 of nine tags, containing all "one" bits for illustrative purposes. It will be noted that the perspective targets are designed to be shared between adjacent tags. Figure 6 shows a square tiling of 16 tags and a corresponding minimum field of view 193, which spans the diagonals of two tags.

Using a (15, 7) Reed-Solomon code, 112 bits of tag data are redundantly encoded to produce 240 encoded bits. The four codewords are interleaved spatially within the tag to

maximize resilience to burst errors. Assuming a 16-bit tag ID as before, this allows a region ID of up to 92 bits.

The data-bearing macrodots 748 of the tag are designed to not overlap their neighbors, so that groups of tags cannot produce structures that resemble targets. This also saves ink. The perspective targets allow detection of the tag, so further targets are not required.

Although the tag may contain an orientation feature to allow disambiguation of the four possible orientations of the tag relative to the sensor, the present invention is concerned with embedding orientation data in the tag data. For example, the four codewords can be arranged so that each tag orientation (in a rotational sense) contains one codeword placed at that orientation, as shown in Figure 5a, where each symbol is labelled with the number of its codeword (1-4) and the position of the symbol within the codeword (A-O). Tag decoding then consists of decoding one codeword at each rotational orientation. Each codeword can either contain a single bit indicating whether it is the first codeword, or two bits indicating which codeword it is. The latter approach has the advantage that if, say, the data content of only one codeword is required, then at most two codewords need to be decoded to obtain the desired data. This may be the case if the region ID is not expected to change within a stroke and is thus only decoded at the start of a stroke. Within a stroke only the codeword containing the tag ID is then desired. Furthermore, since the rotation of the sensing device changes slowly and predictably within a stroke, only one codeword typically needs to be decoded per frame.

It is possible to dispense with perspective targets altogether and instead rely on the data representation being self-registering. In this case each bit value (or multi-bit value) is typically represented by an explicit glyph, i.e. no bit value is represented by the absence of a glyph. This ensures that the data grid is well-populated, and thus allows the grid to be reliably identified and its perspective distortion detected and subsequently corrected during data sampling. To allow tag boundaries to be detected, each tag data must contain a marker pattern, and these must be redundantly encoded to allow reliable detection. The overhead of such marker patterns is similar to the overhead of explicit perspective targets. Various such schemes are described in the present applicants' co-pending PCT application PCT/AU01/01274 filed 11 October 2001.

The arrangement 728 of Figure 5c shows that the square tag 726 can be used to fully tile or tessellate, i.e. without gaps or overlap, a plane of arbitrary size.

Although in preferred embodiments the tagging schemes described herein encode a single data bit using the presence or absence of a single undifferentiated macrodot, they can also use sets of differentiated glyphs to represent single-bit or multi-bit values, such as

the sets of glyphs illustrated in the present applicants' co-pending PCT application PCT/AU01/01274 filed 11 October 2001.

1.2.4.1 Macrodot Packing Schemes

Figure 44 shows a triangular macrodot packing 700 with a four-bit symbol unit 702 outlined. The area of the symbol unit is given by $A_{UNIT} = 2\sqrt{3}s^2 \cong 3.5s^2$, where s the spacing of adjacent macrodots. Figure 45 shows a square macrodot packing 704 with a four-bit symbol unit 706 outlined. The area of the symbol unit is given by $A_{UNIT} = 4s^2$. Figure 46 shows a hexagonal macrodot packing 708 with a four-bit symbol unit 710 outlined. The area of the symbol unit is given by $A_{UNIT} = 3\sqrt{3}s^2 \cong 5.2s^2$. Of these packing schemes, the triangular packing scheme gives the greatest macrodot density for a particular macrodot spacing s .

In preferred embodiments, s has a value between 100 μm and 200 μm .

1.2.4.2 Tag Designs

Figure 46 shows a one-sixth segment 712 of a hexagonal tag, with the segment containing a maximum of 11 four-bit symbols with the triangular macrodot packing shown in Figure 44. The target 17 is shared with adjacent segments. Each tag segment can, by way of example, support a codeword of an $(11, k)$ Reed-Solomon code, i.e. a punctured $(15, k)$ code, with the ability to detect $u = 11 - k$ symbol errors, or correct $t = \lfloor (11 - k)/2 \rfloor$ symbol errors. For example, if $k = 7$ then $u = 4$ and $t = 2$.

Figure 47 shows a one-quarter segment 718 of a square tag, with the segment containing a maximum of 15 four-bit symbols with the square macrodot packing shown in Figure 45. Each tag segment can, by way of example, support a codeword of a $(15, k)$ Reed-Solomon code, with the ability to detect $u = 15 - k$ symbol errors, or correct $t = \lfloor (15 - k)/2 \rfloor$ symbol errors. For example, if $k = 7$ then $u = 8$ and $t = 4$.

1.2.4.3 Hexagonal Tag Design

Figure 48 shows a logical layout of a hexagonal tag 722 using the tag segment 712 of Figure 46, with six interleaved 2^4 -ary $(11, k)$ codewords. Figure 49 shows the macrodot layout of the hexagonal tag 722 of Figure 51. Figure 53 shows an arrangement 724 of seven abutting tags 722 of the design of Figure 48, with shared targets 17. The arrangement 724 shows that the hexagonal tag 722 can be used to tessellate a plane of arbitrary size.

1.2.4.4 Alternative Hexagonal Tag Design 1

Figure 51 shows the logical layout of an alternative hexagonal tag. This tag design is described in detail in the present applicants' co-filed US application USSN _____ entitled "Orientation-Indicating Cyclic Position Codes" (docket number NPT019US).

The tag contains a 2^4 -ary (6, 1) cyclic position codeword (0, 5, 6, 9, A₁₆, F₁₆) which can be decoded at any of the six possible orientations of the tag to determine the actual orientation of the tag. Symbols which are part of the cyclic position codeword have a prefix of "R" and are numbered 0 to 5 in order of increasing significance, and are shown shaded in Figure 52.

The tag locally contains three complete codewords which are used to encode information unique to the tag. Each codeword is of a punctured 2^4 -ary (9, 5) Reed-Solomon code. The tag therefore encodes up to 60 bits of information unique to the tag. The tag also contains fragments of three codewords which are distributed across three adjacent tags and which are used to encode information common to a set of contiguous tags. Each codeword is of a punctured 2^4 -ary (9, 5) Reed-Solomon code. Any three adjacent tags therefore together encode up to 60 bits of information common to a set of contiguous tags.

The layout of the three complete codewords, distributed across three adjacent tags, is shown in Figure 53. In relation to these distributed codewords there are three types of tag. These are referred to as P, Q and R in order of increasing significance.

The P, Q and R tags are repeated in a continuous tiling of tags which guarantees that any set of three adjacent tags contains one tag of each type, and therefore contains a complete set of distributed codewords. The tag type, used to determine the registration of the distributed codewords with respect to a particular set of adjacent tags, is encoded in one of the local codewords of each tag.

1.2.4.4 Alternative Hexagonal Tag Design 2

Figure 54 shows the logical layout of another alternative hexagonal tag. This tag design is described in detail in the present applicants' co-filed US application USSN _____ entitled "Symmetric Tags" (docket number NPT023US).

Figure 54 shows a logical layout of a hexagonal tag 750 using the tag segment of Figure 46, with one local 2^4 -ary (12, k) codeword interleaved with eighteen 3-symbol fragments of eighteen distributed 2^4 -ary (9, k) codewords.

In the layout of Figure 54, the twelve 4-bit symbols of the local codeword are labelled G1 through G12, and are shown with a dashed outline. Each symbol of the eighteen fragments of the eighteen distributed codewords is labelled with an initial prefix of A through F, indicating which of six nominal codewords the symbol belongs to, a subsequent prefix of S through U, indicating which 3-symbol part of the codeword the symbol belongs to, and a suffix of 1 through 3, indicating which of the three possible symbols the symbol is.

Tag 750 is structured so that the minimal field of view allows the recovery of the local codeword G of at least one tag, and the entire set of distributed codewords AP through FR via fragments of tags of type P, Q and R included in the field of view. Furthermore, the continuous tiling of tag 750 ensures that there is a codeword available with a known layout for each possible rotational and translational combination (of which there are eighteen). Each distributed codeword includes data which identifies the rotation of the codeword in relation to the tiling, thus allowing the rotation of the tiling with respect to the field of view to be determined from decoded data rather than from other structures, and the local codeword to be decoded at the correct orientation.

Figure 55 shows the logical layout of the hexagonal tag 750 of Figure 54, re-arranged to show the distributed 3-symbol fragments which contribute to the same codewords. For example, if the central tag shown in Figure 54 were a P-type tag, then the six distributed codewords shown in the figure would be the AP, BP, CP, DP, EP and FP codewords. Figure 55 also shows the local G codeword of the tag. Clearly, given the distributed and repeating nature of the distributed codewords, different fragments from the ones shown in the figure can be used to build the corresponding codewords.

1.2.4 Tag Image Processing and Decoding

Figure 7 shows a tag image processing and decoding process flow. A raw image 202 of the tag pattern is acquired (at 200), for example via an image sensor such as a CCD image sensor, CMOS image sensor, or a scanning laser and photodiode image sensor. The raw image is then typically enhanced (at 204) to produce an enhanced image 206 with improved contrast and more uniform pixel intensities. Image enhancement may include global or local range expansion, equalisation, and the like. The enhanced image 206 is then typically filtered (at 208) to produce a filtered image 210. Image filtering may consist of low-pass filtering, with the low-pass filter kernel size tuned to obscure macrodots but to preserve targets. The filtering step 208 may include additional filtering (such as edge detection) to enhance target features. The filtered image 210 is then processed to locate target features (at 212), yielding a set of target points. This may consist of a search for target features whose spatial inter-relationship is consistent with the known geometry of a tag. Candi-

date targets may be identified directly from maxima in the filtered image 210, or may be the subject of further characterisation and matching, such as via their (binary or grayscale) shape moments (typically computed from pixels in the enhanced image 206 based on local maxima in the filtered image 210), as described in US patent application serial number 09/575,154. The search typically starts from the center of the field of view. The target points 214 found by the search step 212 indirectly identify the location of the tag in the three-dimensional space occupied by the image sensor and its associated optics. Since the target points 214 are derived from the (binary or grayscale) centroids of the targets, they are typically defined to sub-pixel precision.

It may be useful to determine the actual 3D transform of the tag (at 216), and, by extension, the 3D transform (or pose) 218 of the sensing device relative to the tag. This may be done analytically, as described in US patent application serial number 09/575,154, or using a maximum likelihood estimator (such as least squares adjustment) to fit parameter values to the 3D transform given the observed perspective-distorted target points (as described in P.R. Wolf and B.A. Dewitt, *Elements of Photogrammetry with Applications in GIS*, 3rd Edition, McGraw Hill, February 2000, the contents of which are herein incorporated by reference thereto). The 3D transform includes the 3D translation of the tag, the 3D orientation (rotation) of the tag, and the focal length and viewport scale of the sensing device, thus giving eight parameters to be fitted, or six parameters if the focal length and viewport scale are known (e.g. by design or from a calibration step). Each target point yields a pair of observation equations, relating an observed coordinate to a known coordinate. If eight parameters are being fitted, then five or more target points are needed to provide sufficient redundancy to allow maximum likelihood estimation. If six parameters are being fitted, then four or more target points are needed. If the tag design contains more targets than are minimally required to allow maximum likelihood estimation, then the tag can be recognised and decoded even if up to that many of its targets are damaged beyond recognition.

To allow macrodot values to be sampled accurately, the perspective transform of the tag must be inferred. Four of the target points are taken to be the perspective-distorted corners of a rectangle of known size in tag space, and the eight-degree-of-freedom perspective transform 222 is inferred (at 220), based on solving the well-understood equations relating the four tag-space and image-space point pairs (see Heckbert, P., *Fundamentals of Texture Mapping and Image Warping*, Masters Thesis, Dept. of EECS, U. of California at Berkeley, Technical Report No. UCB/CSD 89/516, June 1989, the contents of which are herein incorporated by reference thereto). The perspective transform may alternatively be derived from the 3D transform 218, if available.

The inferred tag-space to image-space perspective transform 222 is used to project (at 224) each known data bit position in tag space into image space where the real-valued position is used to bi-linearly (or higher-order) interpolate (at 224) the four (or more) relevant adjacent pixels in the enhanced input image 206. The resultant macrodot value is compared with a suitable threshold to determine whether it represents a zero bit or a one bit.

Once the bits of one or more complete codeword have been sampled, the codewords are decoded (at 228) to obtain the desired data 230 encoded in the tag. Redundancy in the codeword may be used to detect errors in the sampled data, or to correct errors in the sampled data.

As discussed in US patent application serial number 09/575,154, the obtained tag data 230 may directly or indirectly identify the surface region containing the tag and the position of the tag within the region. An accurate position of the sensing device relative to the surface region can therefore be derived from the tag data 230 and the 3D transform 218 of the sensing device relative to the tag.

1.2.6 Tag Map

Decoding a tag results in a region ID, a tag ID, and a tag-relative pen transform. Before the tag ID and the tag-relative pen location can be translated into an absolute location within the tagged region, the location of the tag within the region must be known. This is given by a tag map, a function which maps each tag ID in a tagged region to a corresponding location. The tag map class diagram is shown in Figure 22, as part of the netpage printer class diagram.

A tag map reflects the scheme used to tile the surface region with tags, and this can vary according to surface type. When multiple tagged regions share the same tiling scheme and the same tag numbering scheme, they can also share the same tag map.

The tag map for a region must be retrievable via the region ID. Thus, given a region ID, a tag ID and a pen transform, the tag map can be retrieved, the tag ID can be translated into an absolute tag location within the region, and the tag-relative pen location can be added to the tag location to yield an absolute pen location within the region.

The tag ID may have a structure which assists translation through the tag map. It may, for example, encode cartesian coordinates or polar coordinates, depending on the surface type on which it appears. The tag ID structure is dictated by and known to the tag map, and tag IDs associated with different tag maps may therefore have different structures. For example, the tag ID may simply encode a pair of x and y coordinates of the tag, in which

case the tag map may simply consist of record of the coordinate precision. If the coordinate precision is fixed, then the tag map can be implicit.

1.2.7 Tagging Schemes

Two distinct surface coding schemes are of interest, both of which use the tag structure described earlier in this section. The preferred coding scheme uses "location-indicating" tags as already discussed. An alternative coding scheme uses object-indicating tags.

A location-indicating tag contains a tag ID which, when translated through the tag map associated with the tagged region, yields a unique tag location within the region. The tag-relative location of the pen is added to this tag location to yield the location of the pen within the region. This in turn is used to determine the location of the pen relative to a user interface element in the page description associated with the region. Not only is the user interface element itself identified, but a location relative to the user interface element is identified. Location-indicating tags therefore trivially support the capture of an absolute pen path in the zone of a particular user interface element.

An object-indicating tag contains a tag ID which directly identifies a user interface element in the page description associated with the region. All the tags in the zone of the user interface element identify the user interface element, making them all identical and therefore indistinguishable. Object-indicating tags do not, therefore, support the capture of an absolute pen path. They do, however, support the capture of a relative pen path. So long as the position sampling frequency exceeds twice the encountered tag frequency, the displacement from one sampled pen position to the next within a stroke can be unambiguously determined.

With either tagging scheme, the tags function in cooperation with associated visual elements on the netpage as user interactive elements in that a user can interact with the printed page using an appropriate sensing device in order for tag data to be read by the sensing device and for an appropriate response to be generated in the netpage system.

1.3 DOCUMENT AND PAGE DESCRIPTIONS

A preferred embodiment of a document and page description class diagram is shown in Figures 25 and 26.

In the netpage system a document is described at three levels. At the most abstract level the document 836 has a hierarchical structure whose terminal elements 839 are associated with content objects 840 such as text objects, text style objects, image objects, etc. Once the document is printed on a printer with a particular page size and according to a par-

ticular user's scale factor preference, the document is paginated and otherwise formatted. Formatted terminal elements 835 will in some cases be associated with content objects which are different from those associated with their corresponding terminal elements, particularly where the content objects are style-related. Each printed instance of a document and page is also described separately, to allow input captured through a particular page instance 830 to be recorded separately from input captured through other instances of the same page description.

The presence of the most abstract document description on the page server allows a user to request a copy of a document without being forced to accept the source document's specific format. The user may be requesting a copy through a printer with a different page size, for example. Conversely, the presence of the formatted document description on the page server allows the page server to efficiently interpret user actions on a particular printed page.

A formatted document 834 consists of a set of formatted page descriptions 5, each of which consists of a set of formatted terminal elements 835. Each formatted element has a spatial extent or zone 58 on the page. This defines the active area of input elements such as hyperlinks and input fields.

A document instance 831 corresponds to a formatted document 834. It consists of a set of page instances 830, each of which corresponds to a page description 5 of the formatted document. Each page instance 830 describes a single unique printed netpage 1, and records the page ID 50 of the netpage. A page instance is not part of a document instance if it represents a copy of a page requested in isolation.

A page instance consists of a set of terminal element instances 832. An element instance only exists if it records instance-specific information. Thus, a hyperlink instance exists for a hyperlink element because it records a transaction ID 55 which is specific to the page instance, and a field instance exists for a field element because it records input specific to the page instance. An element instance does not exist, however, for static elements such as textflows.

A terminal element can be a static element 843, a hyperlink element 844, a field element 845 or a page server command element 846, as shown in Figure 27. A static element 843 can be a style element 847 with an associated style object 854, a textflow element 848 with an associated styled text object 855, an image element 849 with an associated image object 856, a graphic element 850 with an associated graphic object 857, a video clip element 851 with an associated video clip object 858, an audio clip element 852 with an

associated audio clip object 859, or a script element 853 with an associated script object 860, as shown in Figure 28.

A page instance has a background field 833 which is used to record any digital ink captured on the page which does not apply to a specific input element.

In the preferred form of the invention, a tag map 811 is associated with each page instance to allow tags on the page to be translated into locations on the page.

1.4 THE NETPAGE NETWORK

In a preferred embodiment, a netpage network consists of a distributed set of netpage page servers 10, netpage registration servers 11, netpage ID servers 12, netpage application servers 13, netpage publication servers 14, Web terminals 75, netpage printers 601, and relay devices 44 connected via a network 19 such as the Internet, as shown in Figure 3.

The netpage registration server 11 is a server which records relationships between users, pens, printers, applications and publications, and thereby authorizes various network activities. It authenticates users and acts as a signing proxy on behalf of authenticated users in application transactions. It also provides handwriting recognition services. As described above, a netpage page server 10 maintains persistent information about page descriptions and page instances. The netpage network includes any number of page servers, each handling a subset of page instances. Since a page server also maintains user input values for each page instance, clients such as netpage printers send netpage input directly to the appropriate page server. The page server interprets any such input relative to the description of the corresponding page.

A netpage ID server 12 allocates document IDs 51 on demand, and provides load-balancing of page servers via its ID allocation scheme.

A netpage printer uses the Internet Distributed Name System (DNS), or similar, to resolve a netpage page ID 50 into the network address of the netpage page server handling the corresponding page instance.

A netpage application server 13 is a server which hosts interactive netpage applications. A netpage publication server 14 is an application server which publishes netpage documents to netpage printers. They are described in detail in Section 2.

Netpage servers can be hosted on a variety of network server platforms from manufacturers such as IBM, Hewlett-Packard, and Sun. Multiple netpage servers can run concurrently on a single host, and a single server can be distributed over a number of hosts.

Some or all of the functionality provided by netpage servers, and in particular the functionality provided by the ID server and the page server, can also be provided directly in a netpage appliance such as a netpage printer, in a computer workstation, or on a local network.

1.5 THE NETPAGE PRINTER

The netpage printer 601 is an appliance which is registered with the netpage system and prints netpage documents on demand and via subscription. Each printer has a unique printer ID 62, and is connected to the netpage network via a network such as the Internet, ideally via a broadband connection.

Apart from identity and security settings in non-volatile memory, the netpage printer contains no persistent storage. As far as a user is concerned, "the network is the computer". Netpages function interactively across space and time with the help of the distributed netpage page servers 10, independently of particular netpage printers.

The netpage printer receives subscribed netpage documents from netpage publication servers 14. Each document is distributed in two parts: the page layouts, and the actual text and image objects which populate the pages. Because of personalization, page layouts are typically specific to a particular subscriber and so are pointcast to the subscriber's printer via the appropriate page server. Text and image objects, on the other hand, are typically shared with other subscribers, and so are multicast to all subscribers' printers and the appropriate page servers.

The netpage publication server optimizes the segmentation of document content into pointcasts and multicasts. After receiving the pointcast of a document's page layouts, the printer knows which multicasts, if any, to listen to.

Once the printer has received the complete page layouts and objects that define the document to be printed, it can print the document.

The printer rasterizes and prints odd and even pages simultaneously on both sides of the sheet. It contains duplexed print engine controllers 760 and print engines utilizing Memjet™ printheads 350 for this purpose.

The printing process consists of two decoupled stages: rasterization of page descriptions, and expansion and printing of page images. The raster image processor (RIP) consists of one or more standard DSPs 757 running in parallel. The duplexed print engine controllers consist of custom processors which expand, dither and print page images in real time, synchronized with the operation of the printheads in the print engines.

Printers not enabled for IR printing have the option to print tags using IR-absorptive black ink, although this restricts tags to otherwise empty areas of the page. Although such pages have more limited functionality than IR-printed pages, they are still classed as netpages.

A normal netpage printer prints netpages on sheets of paper. More specialised netpage printers may print onto more specialised surfaces, such as globes. Each printer supports at least one surface type, and supports at least one tag tiling scheme, and hence tag map, for each surface type. The tag map 811 which describes the tag tiling scheme actually used to print a document becomes associated with that document so that the document's tags can be correctly interpreted.

Figure 2 shows the netpage printer class diagram, reflecting printer-related information maintained by a registration server 11 on the netpage network.

A preferred embodiment of the netpage printer is described in greater detail in Section 6 below, with reference to Figures 11 to 16.

1.5.1 Memjet™ Printheads

The netpage system can operate using printers made with a wide range of digital printing technologies, including thermal inkjet, piezoelectric inkjet, laser electrophotographic, and others. However, for wide consumer acceptance, it is desirable that a netpage printer have the following characteristics:

- photographic quality color printing
- high quality text printing
- high reliability
- low printer cost
- low ink cost
- low paper cost
- simple operation
- nearly silent printing
- high printing speed
- simultaneous double sided printing
- compact form factor
- low power consumption

No commercially available printing technology has all of these characteristics.

To enable the production of printers with these characteristics, the present applicant has invented a new print technology, referred to as Memjet™ technology. Memjet™ is a drop-on-demand inkjet technology that incorporates pagewidth printheads fabricated using

microelectromechanical systems (MEMS) technology. Figure 17 shows a single printing element 300 of a Memjet™ printhead. The netpage wallprinter incorporates 168960 printing elements 300 to form a 1600 dpi pagewidth duplex printer. This printer simultaneously prints cyan, magenta, yellow, black, and infrared inks as well as paper conditioner and ink fixative.

The printing element 300 is approximately 110 microns long by 32 microns wide. Arrays of these printing elements are formed on a silicon substrate 301 that incorporates CMOS logic, data transfer, timing, and drive circuits (not shown).

Major elements of the printing element 300 are the nozzle 302, the nozzle rim 303, the nozzle chamber 304, the fluidic seal 305, the ink channel rim 306, the lever arm 307, the active actuator beam pair 308, the passive actuator beam pair 309, the active actuator anchor 310, the passive actuator anchor 311, and the ink inlet 312.

The active actuator beam pair 308 is mechanically joined to the passive actuator beam pair 309 at the join 319. Both beams pairs are anchored at their respective anchor points 310 and 311. The combination of elements 308, 309, 310, 311, and 319 form a cantilevered electrothermal bend actuator 320.

While printing, the printhead CMOS circuitry distributes data from the print engine controller to the correct printing element, latches the data, and buffers the data to drive the electrodes 318 of the active actuator beam pair 308. This causes an electrical current to pass through the beam pair 308 for about one microsecond, resulting in Joule heating. The temperature increase resulting from Joule heating causes the beam pair 308 to expand. As the passive actuator beam pair 309 is not heated, it does not expand, resulting in a stress difference between the two beam pairs. This stress difference is partially resolved by the cantilevered end of the electrothermal bend actuator 320 bending towards the substrate 301. The lever arm 307 transmits this movement to the nozzle chamber 304. The nozzle chamber 304 moves about two microns to the position shown in Figure 19(b). This increases the ink pressure, forcing ink 321 out of the nozzle 302, and causing the ink meniscus 316 to bulge. The nozzle rim 303 prevents the ink meniscus 316 from spreading across the surface of the nozzle chamber 304.

As the temperature of the beam pairs 308 and 309 equalizes, the actuator 320 returns to its original position. This aids in the break-off of the ink droplet 317 from the ink 321 in the nozzle chamber. The nozzle chamber is refilled by the action of the surface tension at the meniscus 316.

In a netpage printer, the length of the printhead is the full width of the paper (typically 210 mm). When printing, the paper is moved past the fixed printhead. The printhead

has 6 rows of interdigitated printing elements 300, printing the six colors or types of ink supplied by the ink inlets.

To protect the fragile surface of the printhead during operation, a nozzle guard wafer is attached to the printhead substrate. For each nozzle there is a corresponding nozzle guard hole through which the ink droplets are fired. To prevent the nozzle guard holes from becoming blocked by paper fibers or other debris, filtered air is pumped through the air inlets and out of the nozzle guard holes during printing. To prevent ink from drying, the nozzle guard is sealed while the printer is idle.

1.6 THE NETPAGE PEN

The active sensing device of the netpage system is typically a pen 101, which, using its embedded controller 134, is able to capture and decode IR position tags from a page via an image sensor. The image sensor is a solid-state device provided with an appropriate filter to permit sensing at only near-infrared wavelengths. As described in more detail below, the system is able to sense when the nib is in contact with the surface, and the pen is able to sense tags at a sufficient rate to capture human handwriting (i.e. at 200 dpi or greater and 100 Hz or faster). Information captured by the pen is encrypted and wirelessly transmitted to the printer (or base station), the printer or base station interpreting the data with respect to the (known) page structure.

The preferred embodiment of the netpage pen operates both as a normal marking ink pen and as a non-marking stylus. The marking aspect, however, is not necessary for using the netpage system as a browsing system, such as when it is used as an Internet interface. Each netpage pen is registered with the netpage system and has a unique pen ID 61. Figure 23 shows the netpage pen class diagram, reflecting pen-related information maintained by a registration server 11 on the netpage network.

When either nib is in contact with a netpage, the pen determines its position and orientation relative to the page. The nib is attached to a force sensor, and the force on the nib is interpreted relative to a threshold to indicate whether the pen is "up" or "down". This allows a interactive element on the page to be 'clicked' by pressing with the pen nib, in order to request, say, information from a network. Furthermore, the force is captured as a continuous value to allow, say, the full dynamics of a signature to be verified.

The pen determines the position and orientation of its nib on the netpage by imaging, in the infrared spectrum, an area 193 of the page in the vicinity of the nib. It decodes the nearest tag and computes the position of the nib relative to the tag from the observed perspective distortion on the imaged tag and the known geometry of the pen optics. Although

the position resolution of the tag may be low, because the tag density on the page is inversely proportional to the tag size, the adjusted position resolution is quite high, exceeding the minimum resolution required for accurate handwriting recognition.

Pen actions relative to a netpage are captured as a series of strokes. A stroke consists of a sequence of time-stamped pen positions on the page, initiated by a pen-down event and completed by the subsequent pen-up event. A stroke is also tagged with the page ID 50 of the netpage whenever the page ID changes, which, under normal circumstances, is at the commencement of the stroke.

Each netpage pen has a current selection 826 associated with it, allowing the user to perform copy and paste operations etc. The selection is timestamped to allow the system to discard it after a defined time period. The current selection describes a region of a page instance. It consists of the most recent digital ink stroke captured through the pen relative to the background area of the page. It is interpreted in an application-specific manner once it is submitted to an application via a selection hyperlink activation.

Each pen has a current nib 824. This is the nib last notified by the pen to the system. In the case of the default netpage pen described above, either the marking black ink nib or the non-marking stylus nib is current. Each pen also has a current nib style 825. This is the nib style last associated with the pen by an application, e.g. in response to the user selecting a color from a palette. The default nib style is the nib style associated with the current nib. Strokes captured through a pen are tagged with the current nib style. When the strokes are subsequently reproduced, they are reproduced in the nib style with which they are tagged.

Whenever the pen is within range of a printer with which it can communicate, the pen slowly flashes its "online" LED. When the pen fails to decode a stroke relative to the page, it momentarily activates its "error" LED. When the pen succeeds in decoding a stroke relative to the page, it momentarily activates its "ok" LED.

A sequence of captured strokes is referred to as digital ink. Digital ink forms the basis for the digital exchange of drawings and handwriting, for online recognition of handwriting, and for online verification of signatures.

The pen is wireless and transmits digital ink to the netpage printer via a short-range radio link. The transmitted digital ink is encrypted for privacy and security and packetized for efficient transmission, but is always flushed on a pen-up event to ensure timely handling in the printer.

When the pen is out-of-range of a printer it buffers digital ink in internal memory, which has a capacity of over ten minutes of continuous handwriting. When the pen is once again within range of a printer, it transfers any buffered digital ink.

A pen can be registered with any number of printers, but because all state data resides in netpages both on paper and on the network, it is largely immaterial which printer a pen is communicating with at any particular time.

A preferred embodiment of the pen is described in greater detail in Section 6 below, with reference to Figures 8 to 10.

1.7 NETPAGE INTERACTION

The netpage printer 601 receives data relating to a stroke from the pen 101 when the pen is used to interact with a netpage 1. The coded data 3 of the tags 4 is read by the pen when it is used to execute a movement, such as a stroke. The data allows the identity of the particular page and associated interactive element to be determined and an indication of the relative positioning of the pen relative to the page to be obtained. The indicating data is transmitted to the printer, where it resolves, via the DNS, the page ID 50 of the stroke into the network address of the netpage page server 10 which maintains the corresponding page instance 830. It then transmits the stroke to the page server. If the page was recently identified in an earlier stroke, then the printer may already have the address of the relevant page server in its cache. Each netpage consists of a compact page layout maintained persistently by a netpage page server (see below). The page layout refers to objects such as images, fonts and pieces of text, typically stored elsewhere on the netpage network.

When the page server receives the stroke from the pen, it retrieves the page description to which the stroke applies, and determines which element of the page description the stroke intersects. It is then able to interpret the stroke in the context of the type of the relevant element.

A "click" is a stroke where the distance and time between the pen down position and the subsequent pen up position are both less than some small maximum. An object which is activated by a click typically requires a click to be activated, and accordingly, a longer stroke is ignored. The failure of a pen action, such as a "sloppy" click, to register is indicated by the lack of response from the pen's "ok" LED.

There are two kinds of input elements in a netpage page description: hyperlinks and form fields. Input through a form field can also trigger the activation of an associated hyperlink.

1.7.1 Hyperlinks

A hyperlink is a means of sending a message to a remote application, and typically elicits a printed response in the netpage system.

A hyperlink element 844 identifies the application 71 which handles activation of the hyperlink, a link ID 54 which identifies the hyperlink to the application, an "alias required" flag which asks the system to include the user's application alias ID 65 in the hyperlink activation, and a description which is used when the hyperlink is recorded as a favorite or appears in the user's history. The hyperlink element class diagram is shown in Figure 29.

When a hyperlink is activated, the page server sends a request to an application somewhere on the network. The application is identified by an application ID 64, and the application ID is resolved in the normal way via the DNS. There are three types of hyperlinks: general hyperlinks 863, form hyperlinks 865, and selection hyperlinks 864, as shown in Figure 30. A general hyperlink can implement a request for a linked document, or may simply signal a preference to a server. A form hyperlink submits the corresponding form to the application. A selection hyperlink submits the current selection to the application. If the current selection contains a single-word piece of text, for example, the application may return a single-page document giving the word's meaning within the context in which it appears, or a translation into a different language. Each hyperlink type is characterized by what information is submitted to the application.

The corresponding hyperlink instance 862 records a transaction ID 55 which can be specific to the page instance on which the hyperlink instance appears. The transaction ID can identify user-specific data to the application, for example a "shopping cart" of pending purchases maintained by a purchasing application on behalf of the user.

The system includes the pen's current selection 826 in a selection hyperlink activation. The system includes the content of the associated form instance 868 in a form hyperlink activation, although if the hyperlink has its "submit delta" attribute set, only input since the last form submission is included. The system includes an effective return path in all hyperlink activations.

A hyperlinked group 866 is a group element 838 which has an associated hyperlink, as shown in Figure 31. When input occurs through any field element in the group, the hyperlink 844 associated with the group is activated. A hyperlinked group can be used to associate hyperlink behavior with a field such as a checkbox. It can also be used, in conjunction with the "submit delta" attribute of a form hyperlink, to provide continuous input

to an application. It can therefore be used to support a "blackboard" interaction model, i.e. where input is captured and therefore shared as soon as it occurs.

1.7.2 Forms

A form defines a collection of related input fields used to capture a related set of inputs through a printed netpage. A form allows a user to submit one or more parameters to an application software program running on a server.

A form 867 is a group element 838 in the document hierarchy. It ultimately contains a set of terminal field elements 839. A form instance 868 represents a printed instance of a form. It consists of a set of field instances 870 which correspond to the field elements 845 of the form. Each field instance has an associated value 871, whose type depends on the type of the corresponding field element. Each field value records input through a particular printed form instance, i.e. through one or more printed netpages. The form class diagram is shown in Figure 32.

Each form instance has a status 872 which indicates whether the form is active, frozen, submitted, void or expired. A form is active when first printed. A form becomes frozen once it is signed or once its freeze time is reached. A form becomes submitted once one of its submission hyperlinks has been activated, unless the hyperlink has its "submit delta" attribute set. A form becomes void when the user invokes a void form, reset form or duplicate form page command. A form expires when its specified expiry time is reached, i.e. when the time the form has been active exceeds the form's specified lifetime. While the form is active, form input is allowed. Input through a form which is not active is instead captured in the background field 833 of the relevant page instance. When the form is active or frozen, form submission is allowed. Any attempt to submit a form when the form is not active or frozen is rejected, and instead elicits an form status report.

Each form instance is associated (at 59) with any form instances derived from it, thus providing a version history. This allows all but the latest version of a form in a particular time period to be excluded from a search.

All input is captured as digital ink. Digital ink 873 consists of a set of timestamped stroke groups 874, each of which consists of a set of styled strokes 875. Each stroke consists of a set of timestamped pen positions 876, each of which also includes pen orientation and nib force. The digital ink class diagram is shown in Figure 33.

A field element 845 can be a checkbox field 877, a text field 878, a drawing field 879, or a signature field 880. The field element class diagram is shown in Figure 34. Any digital ink captured in a field's zone 58 is assigned to the field.

A checkbox field has an associated boolean value 881, as shown in Figure 35. Any mark (a tick, a cross, a stroke, a fill zigzag, etc.) captured in a checkbox field's zone causes a true value to be assigned to the field's value.

A text field has an associated text value 882, as shown in Figure 36. Any digital ink captured in a text field's zone is automatically converted to text via online handwriting recognition, and the text is assigned to the field's value. Online handwriting recognition is well-understood (see, for example, Tappert, C., C.Y. Suen and T. Wakahara, "The State of the Art in On-Line Handwriting Recognition", IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol.12, No.8, August 1990, the contents of which are herein incorporated by cross-reference).

A signature field has an associated digital signature value 883, as shown in Figure 37. Any digital ink captured in a signature field's zone is automatically verified with respect to the identity of the owner of the pen, and a digital signature of the content of the form of which the field is part is generated and assigned to the field's value. The digital signature is generated using the pen user's private signature key specific to the application which owns the form. Online signature verification is well-understood (see, for example, Plamondon, R. and G. Lorette, "Automatic Signature Verification and Writer Identification – The State of the Art", Pattern Recognition, Vol.22, No.2, 1989, the contents of which are herein incorporated by cross-reference).

A field element is hidden if its "hidden" attribute is set. A hidden field element does not have an input zone on a page and does not accept input. It can have an associated field value which is included in the form data when the form containing the field is submitted.

"Editing" commands, such as strike-throughs indicating deletion, can also be recognized in form fields.

Because the handwriting recognition algorithm works "online" (i.e. with access to the dynamics of the pen movement), rather than "offline" (i.e. with access only to a bitmap of pen markings), it can recognize run-on discretely-written characters with relatively high accuracy, without a writer-dependent training phase. A writer-dependent model of handwriting is automatically generated over time, however, and can be generated up-front if necessary.

Digital ink, as already stated, consists of a sequence of strokes. Any stroke which starts in a particular element's zone is appended to that element's digital ink stream, ready for interpretation. Any stroke not appended to an object's digital ink stream is appended to the background field's digital ink stream.

Digital ink captured in the background field is interpreted as a selection gesture. Circumscription of one or more objects is generally interpreted as a selection of the circumscribed objects, although the actual interpretation is application-specific.

Table 2 summarises these various pen interactions with a netpage.

Table 2 - Summary of pen interactions with a netpage

Object	Type	Pen Input	Action
Hyperlink	General	Click	Submit action to application
	Form	Click	Submit form to application
	Selection	Click	Submit selection to application
Form field	Checkbox	Any mark	Assign true to field
	Text	Handwriting	Convert digital ink to text; assign text to field
	Drawing	Digital ink	Assign digital ink to field
	Signature	Signature	Verify digital ink signature; generate digital signature of form; assign digital signature to field
None		Circumscription	Assign digital ink to current selection

The system maintains a current selection for each pen. The selection consists simply of the most recent stroke captured in the background field. The selection is cleared after an inactivity timeout to ensure predictable behavior.

The raw digital ink captured in every field is retained on the netpage page server and is optionally transmitted with the form data when the form is submitted to the application. This allows the application to interrogate the raw digital ink should it suspect the original conversion, such as the conversion of handwritten text. This can, for example, involve human intervention at the application level for forms which fail certain application-specific consistency checks. As an extension to this, the entire background area of a form can be designated as a drawing field. The application can then decide, on the basis of the presence of digital ink outside the explicit fields of the form, to route the form to a human operator, on the assumption that the user may have indicated amendments to the filled-in fields outside of those fields.

Figure 38 shows a flowchart of the process of handling pen input relative to a netpage. The process consists of receiving (at 884) a stroke from the pen; identifying (at 885) the page instance 830 to which the page ID 50 in the stroke refers; retrieving (at 886) the page description 5; identifying (at 887) a formatted element 839 whose zone 58 the stroke intersects; determining (at 888) whether the formatted element corresponds to a field

element, and if so appending (at 892) the received stroke to the digital ink of the field value 871, interpreting (at 893) the accumulated digital ink of the field, and determining (at 894) whether the field is part of a hyperlinked group 866 and if so activating (at 895) the associated hyperlink; alternatively determining (at 889) whether the formatted element corresponds to a hyperlink element and if so activating (at 895) the corresponding hyperlink; alternatively, in the absence of an input field or hyperlink, appending (at 890) the received stroke to the digital ink of the background field 833; and copying (at 891) the received stroke to the current selection 826 of the current pen, as maintained by the registration server.

Figure 38a shows a detailed flowchart of step 893 in the process shown in Figure 38, where the accumulated digital ink of a field is interpreted according to the type of the field. The process consists of determining (at 896) whether the field is a checkbox and (at 897) whether the digital ink represents a checkmark, and if so assigning (at 898) a true value to the field value; alternatively determining (at 899) whether the field is a text field and if so converting (at 900) the digital ink to computer text, with the help of the appropriate registration server, and assigning (at 901) the converted computer text to the field value; alternatively determining (at 902) whether the field is a signature field and if so verifying (at 903) the digital ink as the signature of the pen's owner, with the help of the appropriate registration server, creating (at 904) a digital signature of the contents of the corresponding form, also with the help of the registration server and using the pen owner's private signature key relating to the corresponding application, and assigning (at 905) the digital signature to the field value.

1.7.3 Page Server Commands

A page server command is a command which is handled locally by the page server. It operates directly on form, page and document instances.

A page server command 907 can be a void form command 908, a duplicate form command 909, a reset form command 910, a get form status command 911, a duplicate page command 912, a reset page command 913, a get page status command 914, a duplicate document command 915, a reset document command 916, or a get document status command 917, as shown in Figure 39.

A void form command voids the corresponding form instance. A duplicate form command voids the corresponding form instance and then produces an active printed copy of the current form instance with field values preserved. The copy contains the same hyperlink transaction IDs as the original, and so is indistinguishable from the original to an application. A reset form command voids the corresponding form instance and then produces an

active printed copy of the form instance with field values discarded. A get form status command produces a printed report on the status of the corresponding form instance, including who published it, when it was printed, for whom it was printed, and the form status of the form instance.

Since a form hyperlink instance contains a transaction ID, the application has to be involved in producing a new form instance. A button requesting a new form instance is therefore typically implemented as a hyperlink.

A duplicate page command produces a printed copy of the corresponding page instance with the background field value preserved. If the page contains a form or is part of a form, then the duplicate page command is interpreted as a duplicate form command. A reset page command produces a printed copy of the corresponding page instance with the background field value discarded. If the page contains a form or is part of a form, then the reset page command is interpreted as a reset form command. A get page status command produces a printed report on the status of the corresponding page instance, including who published it, when it was printed, for whom it was printed, and the status of any forms it contains or is part of.

The netpage logo which appears on every netpage is usually associated with a duplicate page element.

When a page instance is duplicated with field values preserved, field values are printed in their native form, i.e. a checkmark appears as a standard checkmark graphic, and text appears as typeset text. Only drawings and signatures appear in their original form, with a signature accompanied by a standard graphic indicating successful signature verification.

A duplicate document command produces a printed copy of the corresponding document instance with background field values preserved. If the document contains any forms, then the duplicate document command duplicates the forms in the same way a duplicate form command does. A reset document command produces a printed copy of the corresponding document instance with background field values discarded. If the document contains any forms, then the reset document command resets the forms in the same way a reset form command does. A get document status command produces a printed report on the status of the corresponding document instance, including who published it, when it was printed, for whom it was printed, and the status of any forms it contains.

If the page server command's "on selected" attribute is set, then the command operates on the page identified by the pen's current selection rather than on the page containing the command. This allows a menu of page server commands to be printed. If the tar-

get page doesn't contain a page server command element for the designated page server command, then the command is ignored.

An application can provide application-specific handling by embedding the relevant page server command element in a hyperlinked group. The page server activates the hyperlink associated with the hyperlinked group rather than executing the page server command.

A page server command element is hidden if its "hidden" attribute is set. A hidden command element does not have an input zone on a page and so cannot be activated directly by a user. It can, however, be activated via a page server command embedded in a different page, if that page server command has its "on selected" attribute set.

1.8 STANDARD FEATURES OF NETPAGES

In the preferred form, each netpage is printed with the netpage logo at the bottom to indicate that it is a netpage and therefore has interactive properties. The logo also acts as a copy button. In most cases pressing the logo produces a copy of the page. In the case of a form, the button produces a copy of the entire form. And in the case of a secure document, such as a ticket or coupon, the button elicits an explanatory note or advertising page.

The default single-page copy function is handled directly by the relevant netpage page server. Special copy functions are handled by linking the logo button to an application.

1.9 USER HELP SYSTEM

In a preferred embodiment, the netpage printer has a single button labelled "Help". When pressed it elicits a single help page 46 of information, including:

- status of printer connection
- status of printer consumables
- top-level help menu
- document function menu
- top-level netpage network directory

The help menu provides a hierarchical manual on how to use the netpage system.

The document function menu includes the following functions:

- print a copy of a document
- print a clean copy of a form
- print the status of a document

A document function is initiated by selecting the document and then pressing the button. The status of a document indicates who published it and when, to whom it was delivered, and to whom and when it was subsequently submitted as a form.

The help page is obviously unavailable if the printer is unable to print. In this case the "error" light is lit and the user can request remote diagnosis over the network.

2 PERSONALIZED PUBLICATION MODEL

In the following description, news is used as a canonical publication example to illustrate personalization mechanisms in the netpage system. Although news is often used in the limited sense of newspaper and newsmagazine news, the intended scope in the present context is wider.

In the netpage system, the editorial content and the advertising content of a news publication are personalized using different mechanisms. The editorial content is personalized according to the reader's explicitly stated and implicitly captured interest profile. The advertising content is personalized according to the reader's locality and demographic.

2.1 EDITORIAL PERSONALIZATION

A subscriber can draw on two kinds of news sources: those that deliver news publications, and those that deliver news streams. While news publications are aggregated and edited by the publisher, news streams are aggregated either by a news publisher or by a specialized news aggregator. News publications typically correspond to traditional newspapers and newsmagazines, while news streams can be many and varied: a "raw" news feed from a news service, a cartoon strip, a freelance writer's column, a friend's bulletin board, or the reader's own e-mail.

The netpage publication server supports the publication of edited news publications as well as the aggregation of multiple news streams. By handling the aggregation and hence the formatting of news streams selected directly by the reader, the server is able to place advertising on pages over which it otherwise has no editorial control.

The subscriber builds a daily newspaper by selecting one or more contributing news publications, and creating a personalized version of each. The resulting daily editions are printed and bound together into a single newspaper. The various members of a household typically express their different interests and tastes by selecting different daily publications and then customizing them.

For each publication, the reader optionally selects specific sections. Some sections appear daily, while others appear weekly. The daily sections available from The New York Times online, for example, include "Page One Plus", "National", "International", "Opinion", "Business", "Arts/Living", "Technology", and "Sports". The set of available sections is specific to a publication, as is the default subset.

The reader can extend the daily newspaper by creating custom sections, each one drawing on any number of news streams. Custom sections might be created for e-mail and friends' announcements ("Personal"), or for monitoring news feeds for specific topics ("Alerts" or "Clippings").

For each section, the reader optionally specifies its size, either qualitatively (e.g. short, medium, or long), or numerically (i.e. as a limit on its number of pages), and the desired proportion of advertising, either qualitatively (e.g. high, normal, low, none), or numerically (i.e. as a percentage).

The reader also optionally expresses a preference for a large number of shorter articles or a small number of longer articles. Each article is ideally written (or edited) in both short and long forms to support this preference.

An article may also be written (or edited) in different versions to match the expected sophistication of the reader, for example to provide children's and adults' versions. The appropriate version is selected according to the reader's age. The reader can specify a "reading age" which takes precedence over their biological age.

The articles which make up each section are selected and prioritized by the editors, and each is assigned a useful lifetime. By default they are delivered to all relevant subscribers, in priority order, subject to space constraints in the subscribers' editions.

In sections where it is appropriate, the reader may optionally enable collaborative filtering. This is then applied to articles which have a sufficiently long lifetime. Each article which qualifies for collaborative filtering is printed with rating buttons at the end of the article. The buttons can provide an easy choice (e.g. "liked" and "disliked"), making it more likely that readers will bother to rate the article.

Articles with high priorities and short lifetimes are therefore effectively considered essential reading by the editors and are delivered to most relevant subscribers.

The reader optionally specifies a serendipity factor, either qualitatively (e.g. do or don't surprise me), or numerically. A high serendipity factor lowers the threshold used for matching during collaborative filtering. A high factor makes it more likely that the corre-

sponding section will be filled to the reader's specified capacity. A different serendipity factor can be specified for different days of the week.

The reader also optionally specifies topics of particular interest within a section, and this modifies the priorities assigned by the editors.

The speed of the reader's Internet connection affects the quality at which images can be delivered. The reader optionally specifies a preference for fewer images or smaller images or both. If the number or size of images is not reduced, then images may be delivered at lower quality (i.e. at lower resolution or with greater compression).

At a global level, the reader specifies how quantities, dates, times and monetary values are localized. This involves specifying whether units are imperial or metric, a local timezone and time format, and a local currency, and whether the localization consist of *in situ* translation or annotation. These preferences are derived from the reader's locality by default.

To reduce reading difficulties caused by poor eyesight, the reader optionally specifies a global preference for a larger presentation. Both text and images are scaled accordingly, and less information is accommodated on each page.

The language in which a news publication is published, and its corresponding text encoding, is a property of the publication and not a preference expressed by the user. However, the netpage system can be configured to provide automatic translation services in various guises.

2.2 ADVERTISING LOCALIZATION AND TARGETING

The personalization of the editorial content directly affects the advertising content, because advertising is typically placed to exploit the editorial context. Travel ads, for example, are more likely to appear in a travel section than elsewhere. The value of the editorial content to an advertiser (and therefore to the publisher) lies in its ability to attract large numbers of readers with the right demographics.

Effective advertising is placed on the basis of locality and demographics. Locality determines proximity to particular services, retailers etc., and particular interests and concerns associated with the local community and environment. Demographics determine general interests and preoccupations as well as likely spending patterns.

A news publisher's most profitable product is advertising "space", a multi-dimensional entity determined by the publication's geographic coverage, the size of its readership, its readership demographics, and the page area available for advertising.

In the netpage system, the netpage publication server computes the approximate multi-dimensional size of a publication's saleable advertising space on a per-section basis, taking into account the publication's geographic coverage, the section's readership, the size of each reader's section edition, each reader's advertising proportion, and each reader's demographic.

In comparison with other media, the netpage system allows the advertising space to be defined in greater detail, and allows smaller pieces of it to be sold separately. It therefore allows it to be sold at closer to its true value.

For example, the same advertising "slot" can be sold in varying proportions to several advertisers, with individual readers' pages randomly receiving the advertisement of one advertiser or another, overall preserving the proportion of space sold to each advertiser.

The netpage system allows advertising to be linked directly to detailed product information and online purchasing. It therefore raises the intrinsic value of the advertising space.

Because personalization and localization are handled automatically by netpage publication servers, an advertising aggregator can provide arbitrarily broad coverage of both geography and demographics. The subsequent disaggregation is efficient because it is automatic. This makes it more cost-effective for publishers to deal with advertising aggregators than to directly capture advertising. Even though the advertising aggregator is taking a proportion of advertising revenue, publishers may find the change profit-neutral because of the greater efficiency of aggregation. The advertising aggregator acts as an intermediary between advertisers and publishers, and may place the same advertisement in multiple publications.

It is worth noting that ad placement in a netpage publication can be more complex than ad placement in the publication's traditional counterpart, because the publication's advertising space is more complex. While ignoring the full complexities of negotiations between advertisers, advertising aggregators and publishers, the preferred form of the netpage system provides some automated support for these negotiations, including support for automated auctions of advertising space. Automation is particularly desirable for the placement of advertisements which generate small amounts of income, such as small or highly localized advertisements.

Once placement has been negotiated, the aggregator captures and edits the advertisement and records it on a netpage ad server. Correspondingly, the publisher records the ad placement on the relevant netpage publication server. When the netpage publication server

lays out each user's personalized publication, it picks the relevant advertisements from the netpage ad server.

2.3 USER PROFILES

2.3.1 Information Filtering

The personalization of news and other publications relies on an assortment of user-specific profile information, including:

- publication customizations
- collaborative filtering vectors
- contact details
- presentation preferences

The customization of a publication is typically publication-specific, and so the customization information is maintained by the relevant netpage publication server.

A collaborative filtering vector consists of the user's ratings of a number of news items. It is used to correlate different users' interests for the purposes of making recommendations. Although there are benefits to maintaining a single collaborative filtering vector independently of any particular publication, there are two reasons why it is more practical to maintain a separate vector for each publication: there is likely to be more overlap between the vectors of subscribers to the same publication than between those of subscribers to different publications; and a publication is likely to want to present its users' collaborative filtering vectors as part of the value of its brand, not to be found elsewhere. Collaborative filtering vectors are therefore also maintained by the relevant netpage publication server.

Contact details, including name, street address, ZIP Code, state, country, telephone numbers, are global by nature, and are maintained by a netpage registration server.

Presentation preferences, including those for quantities, dates and times, are likewise global and maintained in the same way.

The localization of advertising relies on the locality indicated in the user's contact details, while the targeting of advertising relies on personal information such as date of birth, gender, marital status, income, profession, education, or qualitative derivatives such as age range and income range.

For those users who choose to reveal personal information for advertising purposes, the information is maintained by the relevant netpage registration server. In the absence of such information, advertising can be targeted on the basis of the demographic associated with the user's ZIP or ZIP+4 Code.

Each user, pen, printer, application provider and application is assigned its own unique identifier, and the netpage registration server maintains the relationships between them, as shown in Figures 21, 22, 23 and 24. For registration purposes, a publisher is a special kind of application provider, and a publication is a special kind of application.

Each user 800 may be authorized to use any number of printers 802, and each printer may allow any number of users to use it. Each user has a single default printer (at 66), to which periodical publications are delivered by default, whilst pages printed on demand are delivered to the printer through which the user is interacting. The server keeps track of which publishers a user has authorized to print to the user's default printer. A publisher does not record the ID of any particular printer, but instead resolves the ID when it is required. The user may also be designated as having administrative privileges 69 on the printer, allowing the user to authorize other users to use the printer. This only has meaning if the printer requires administrative privileges 84 for such operations.

When a user subscribes 808 to a publication 807, the publisher 806 (i.e. application provider 803) is authorized to print to a specified printer or the user's default printer. This authorization can be revoked at any time by the user. Each user may have several pens 801, but a pen is specific to a single user. If a user is authorized to use a particular printer, then that printer recognizes any of the user's pens.

The pen ID is used to locate the corresponding user profile maintained by a particular netpage registration server, via the DNS in the usual way.

A Web terminal 809 can be authorized to print on a particular netpage printer, allowing Web pages and netpage documents encountered during Web browsing to be conveniently printed on the nearest netpage printer.

The netpage system can collect, on behalf of a printer provider, fees and commissions on income earned through publications printed on the provider's printers. Such income can include advertising fees, click-through fees, e-commerce commissions, and transaction fees. If the printer is owned by the user, then the user is the printer provider.

Each user also has a netpage account 820 which is used to accumulate micro-debits and credits (such as those described in the preceding paragraph); contact details 815, including name, address and telephone numbers; global preferences 816, including privacy, delivery and localization settings; any number of biometric records 817, containing the user's encoded signature 818, fingerprint 819 etc; a handwriting model 819 automatically maintained by the system; and SET payment card accounts 821, with which e-commerce payments can be made.

In addition to the user-specific netpage account, each user also has a netpage account 936 specific to each printer the user is authorized to use. Each printer-specific account is used to accumulate micro-debits and credits related to the user's activities on that printer. The user is billed on a regular basis for any outstanding debit balances.

A user optionally appears in the netpage user directory 823, allowing other users to locate and direct e-mail (etc.) to the user.

2.4 INTELLIGENT PAGE LAYOUT

The netpage publication server automatically lays out the pages of each user's personalized publication on a section-by-section basis. Since most advertisements are in the form of pre-formatted rectangles, they are placed on the page before the editorial content.

The advertising ratio for a section can be achieved with wildly varying advertising ratios on individual pages within the section, and the ad layout algorithm exploits this. The algorithm is configured to attempt to co-locate closely tied editorial and advertising content, such as placing ads for roofing material specifically within the publication because of a special feature on do-it-yourself roofing repairs.

The editorial content selected for the user, including text and associated images and graphics, is then laid out according to various aesthetic rules.

The entire process, including the selection of ads and the selection of editorial content, must be iterated once the layout has converged, to attempt to more closely achieve the user's stated section size preference. The section size preference can, however, be matched *on average* over time, allowing significant day-to-day variations.

2.5 DOCUMENT FORMAT

Once the document is laid out, it is encoded for efficient distribution and persistent storage on the netpage network.

The primary efficiency mechanism is the separation of information specific to a single user's edition and information shared between multiple users' editions. The specific information consists of the page layout. The shared information consists of the objects to which the page layout refers, including images, graphics, and pieces of text.

A text object contains fully-formatted text represented in the Extensible Markup Language (XML) using the Extensible Stylesheet Language (XSL). XSL provides precise control over text formatting independently of the region into which the text is being set, which in this case is being provided by the layout. The text object contains embedded lan-

guage codes to enable automatic translation, and embedded hyphenation hints to aid with paragraph formatting.

An image object encodes an image in the JPEG 2000 wavelet-based compressed image format. A graphic object encodes a 2D graphic in Scalable Vector Graphics (SVG) format.

The layout itself consists of a series of placed image and graphic objects, linked textflow objects through which text objects flow, hyperlinks and input fields as described above, and watermark regions. These layout objects are summarized in Table 3. The layout uses a compact format suitable for efficient distribution and storage.

Table 3 - netpage layout objects

Layout object	Attribute	Format of linked object
Image	Position	-
	Image object ID	JPEG 2000
Graphic	Position	-
	Graphic object ID	SVG
Textflow	Textflow ID	-
	Zone	-
	Optional text object ID	XML/XSL
Hyperlink	Type	-
	Zone	-
	Application ID, etc.	-
Field	Type	-
	Meaning	-
	Zone	-
Watermark	Zone	-

2.6 DOCUMENT DISTRIBUTION

As described above, for purposes of efficient distribution and persistent storage on the netpage network, a user-specific page layout is separated from the shared objects to which it refers.

When a subscribed publication is ready to be distributed, the netpage publication server allocates, with the help of the netpage ID server 12, a unique ID for each page, page instance, document, and document instance.

The server computes a set of optimized subsets of the shared content and creates a multicast channel for each subset, and then tags each user-specific layout with the names of the multicast channels which will carry the shared content used by that layout. The server then pointcasts each user's layouts to that user's printer via the appropriate page server, and when the pointcasting is complete, multicasts the shared content on the specified channels. After receiving its pointcast, each page server and printer subscribes to the multicast channels specified in the page layouts. During the multicasts, each page server and printer extracts from the multicast streams those objects referred to by its page layouts. The page servers persistently archive the received page layouts and shared content.

Once a printer has received all the objects to which its page layouts refer, the printer re-creates the fully-populated layout and then rasterizes and prints it.

Under normal circumstances, the printer prints pages faster than they can be delivered. Assuming a quarter of each page is covered with images, the average page has a size of less than 400KB. The printer can therefore hold in excess of 100 such pages in its internal 64MB memory, allowing for temporary buffers etc. The printer prints at a rate of one page per second. This is equivalent to 400KB or about 3Mbit of page data per second, which is similar to the highest expected rate of page data delivery over a broadband network.

Even under abnormal circumstances, such as when the printer runs out of paper, it is likely that the user will be able to replenish the paper supply before the printer's 100-page internal storage capacity is exhausted.

However, if the printer's internal memory does fill up, then the printer will be unable to make use of a multicast when it first occurs. The netpage publication server therefore allows printers to submit requests for re-multicasts. When a critical number of requests is received or a timeout occurs, the server re-multicasts the corresponding shared objects.

Once a document is printed, a printer can produce an exact duplicate at any time by retrieving its page layouts and contents from the relevant page server.

2.7 ON-DEMAND DOCUMENTS

When a netpage document is requested on demand, it can be personalized and delivered in much the same way as a periodical. However, since there is no shared content, delivery is made directly to the requesting printer without the use of multicast.

When a non-netpage document is requested on demand, it is not personalized, and it is delivered via a designated netpage formatting server which reformats it as a netpage document. A netpage formatting server is a special instance of a netpage publication server. The netpage formatting server has knowledge of various Internet document formats, including Adobe's Portable Document Format (PDF), and Hypertext Markup Language (HTML). In the case of HTML, it can make use of the higher resolution of the printed page to present Web pages in a multi-column format, with a table of contents. It can automatically include all Web pages directly linked to the requested page. The user can tune this behavior via a preference.

The netpage formatting server makes standard netpage behavior, including interactivity and persistence, available on any Internet document, no matter what its origin and format. It hides knowledge of different document formats from both the netpage printer and the netpage page server, and hides knowledge of the netpage system from Web servers.

3 SECURITY

3.1 CRYPTOGRAPHY

Cryptography is used to protect sensitive information, both in storage and in transit, and to authenticate parties to a transaction. There are two classes of cryptography in widespread use: secret-key cryptography and public-key cryptography. The netpage network uses both classes of cryptography.

Secret-key cryptography, also referred to as symmetric cryptography, uses the same key to encrypt and decrypt a message. Two parties wishing to exchange messages must first arrange to securely exchange the secret key.

Public-key cryptography, also referred to as asymmetric cryptography, uses two encryption keys. The two keys are mathematically related in such a way that any message encrypted using one key can only be decrypted using the other key. One of these keys is then published, while the other is kept private. The public key is used to encrypt any message intended for the holder of the private key. Once encrypted using the public key, a message can only be decrypted using the private key. Thus two parties can securely exchange messages without first having to exchange a secret key. To ensure that the private key is secure, it is normal for the holder of the private key to generate the key pair.

Public-key cryptography can be used to create a digital signature. The holder of the private key can create a known hash of a message and then encrypt the hash using the private key. Anyone can then verify that the encrypted hash constitutes the "signature" of

the holder of the private key with respect to that particular message by decrypting the encrypted hash using the public key and verifying the hash against the message. If the signature is appended to the message, then the recipient of the message can verify both that the message is genuine and that it has not been altered in transit.

To make public-key cryptography work, there has to be a way to distribute public keys which prevents impersonation. This is normally done using certificates and certificate authorities. A certificate authority is a trusted third party which authenticates the connection between a public key and someone's identity. The certificate authority verifies the person's identity by examining identity documents, and then creates and signs a digital certificate containing the person's identity details and public key. Anyone who trusts the certificate authority can use the public key in the certificate with a high degree of certainty that it is genuine. They just have to verify that the certificate has indeed been signed by the certificate authority, whose public key is well-known.

In most transaction environments, public-key cryptography is only used to create digital signatures and to securely exchange secret session keys. Secret-key cryptography is used for all other purposes.

In the following discussion, when reference is made to the *secure* transmission of information between a netpage printer and a server, what actually happens is that the printer obtains the server's certificate, authenticates it with reference to the certificate authority, uses the public key-exchange key in the certificate to exchange a secret session key with the server, and then uses the secret session key to encrypt the message data. A *session* key, by definition, can have an arbitrarily short lifetime.

3.2 NETPAGE PRINTER SECURITY

Each netpage printer is assigned a pair of unique identifiers at time of manufacture which are stored in read-only memory in the printer and in the netpage registration server database. The first ID 62 is public and uniquely identifies the printer on the netpage network. The second ID is secret and is used when the printer is first registered on the network.

When the printer connects to the netpage network for the first time after installation, it creates a signature public/private key pair. It transmits the secret ID and the public key securely to the netpage registration server. The server compares the secret ID against the printer's secret ID recorded in its database, and accepts the registration if the IDs match. It then creates and signs a certificate containing the printer's public ID and public signature key, and stores the certificate in the registration database.

The netpage registration server acts as a certificate authority for netpage printers, since it has access to secret information allowing it to verify printer identity.

When a user subscribes to a publication, a record is created in the netpage registration server database authorizing the publisher to print the publication to the user's default printer or a specified printer. Every document sent to a printer via a page server is addressed to a particular user and is signed by the publisher using the publisher's private signature key. The page server verifies, via the registration database, that the publisher is authorized to deliver the publication to the specified user. The page server verifies the signature using the publisher's public key, obtained from the publisher's certificate stored in the registration database.

The netpage registration server accepts requests to add printing authorizations to the database, so long as those requests are initiated via a pen registered to the printer.

3.3 NETPAGE PEN SECURITY

Each netpage pen is assigned a unique identifier at time of manufacture which is stored in read-only memory in the pen and in the netpage registration server database. The pen ID 61 uniquely identifies the pen on the netpage network.

A netpage pen can "know" a number of netpage printers, and a printer can "know" a number of pens. A pen communicates with a printer via a radio frequency signal whenever it is within range of the printer. Once a pen and printer are registered, they regularly exchange session keys. Whenever the pen transmits digital ink to the printer, the digital ink is always encrypted using the appropriate session key. Digital ink is never transmitted in the clear.

A pen stores a session key for every printer it knows, indexed by printer ID, and a printer stores a session key for every pen it knows, indexed by pen ID. Both have a large but finite storage capacity for session keys, and will forget a session key on a least-recently-used basis if necessary.

When a pen comes within range of a printer, the pen and printer discover whether they know each other. If they don't know each other, then the printer determines whether it is supposed to know the pen. This might be, for example, because the pen belongs to a user who is registered to use the printer. If the printer is meant to know the pen but doesn't, then it initiates the automatic pen registration procedure. If the printer isn't meant to know the pen, then it agrees with the pen to ignore it until the pen is placed in a charging cup, at which time it initiates the registration procedure.

In addition to its public ID, the pen contains a secret key-exchange key. The key-exchange key is also recorded in the netpage registration server database at time of manufacture. During registration, the pen transmits its pen ID to the printer, and the printer transmits the pen ID to the netpage registration server. The server generates a session key for the printer and pen to use, and securely transmits the session key to the printer. It also transmits a copy of the session key encrypted with the pen's key-exchange key. The printer stores the session key internally, indexed by the pen ID, and transmits the encrypted session key to the pen. The pen stores the session key internally, indexed by the printer ID.

Although a fake pen can impersonate a pen in the pen registration protocol, only a real pen can decrypt the session key transmitted by the printer.

When a previously unregistered pen is first registered, it is of limited use until it is linked to a user. A registered but "un-owned" pen is only allowed to be used to request and fill in netpage user and pen registration forms, to register a new user to which the new pen is automatically linked, or to add a new pen to an existing user.

The pen uses secret-key rather than public-key encryption because of hardware performance constraints in the pen.

3.4 SECURE DOCUMENTS

The netpage system supports the delivery of secure documents such as tickets and coupons. The netpage printer includes a facility to print watermarks, but will only do so on request from publishers who are suitably authorized. The publisher indicates its authority to print watermarks in its certificate, which the printer is able to authenticate.

The "watermark" printing process uses an alternative dither matrix in specified "watermark" regions of the page. Back-to-back pages contain mirror-image watermark regions which coincide when printed. The dither matrices used in odd and even pages' watermark regions are designed to produce an interference effect when the regions are viewed together, achieved by looking *through* the printed sheet.

The effect is similar to a watermark in that it is not visible when looking at only one side of the page, and is lost when the page is copied by normal means.

Pages of secure documents cannot be copied using the built-in netpage copy mechanism described in Section 1.9 above. This extends to copying netpages on netpage-aware photocopiers.

Secure documents are typically generated as part of e-commerce transactions. They can therefore include the user's photograph which was captured when the user registered biometric information with the netpage registration server, as described in Section 2.

When presented with a secure netpage document, the recipient can verify its authenticity by requesting its status in the usual way. The unique ID of a secure document is only valid for the lifetime of the document, and secure document IDs are allocated non-contiguously to prevent their prediction by opportunistic forgers. A secure document verification pen can be developed with built-in feedback on verification failure, to support easy point-of-presentation document verification.

Clearly neither the watermark nor the user's photograph are secure in a cryptographic sense. They simply provide a significant obstacle to casual forgery. Online document verification, particularly using a verification pen, provides an added level of security where it is needed, but is still not entirely immune to forgeries.

3.5 NON-REPUDIATION

In the netpage system, forms submitted by users are delivered reliably to forms handlers and are persistently archived on netpage page servers. It is therefore impossible for recipients to repudiate delivery.

E-commerce payments made through the system, as described in Section 4, are also impossible for the payee to repudiate.

4 ELECTRONIC COMMERCE MODEL

4.1 SECURE ELECTRONIC TRANSACTION (SET)

The netpage system uses the Secure Electronic Transaction (SET) system as one of its payment systems. SET, having been developed by MasterCard and Visa, is organized around payment cards, and this is reflected in the terminology. However, much of the system is independent of the type of accounts being used.

In SET, cardholders and merchants register with a certificate authority and are issued with certificates containing their public signature keys. The certificate authority verifies a cardholder's registration details with the card issuer as appropriate, and verifies a merchant's registration details with the acquirer as appropriate. Cardholders and merchants store their respective private signature keys securely on their computers. During the payment process, these certificates are used to mutually authenticate a merchant and cardholder; and to authenticate them both to the payment gateway.

SET has not yet been adopted widely, partly because cardholder maintenance of keys and certificates is considered burdensome. Interim solutions which maintain cardholder keys and certificates on a server and give the cardholder access via a password have met with some success.

4.2 SET PAYMENTS

In the netpage system the netpage registration server acts as a proxy for the netpage user (i.e. the cardholder) in SET payment transactions.

The netpage system uses biometrics to authenticate the user and authorize SET payments. Because the system is pen-based, the biometric used is the user's on-line signature, consisting of time-varying pen position and pressure. A fingerprint biometric can also be used by designing a fingerprint sensor into the pen, although at a higher cost. The type of biometric used only affects the capture of the biometric, not the authorization aspects of the system.

The first step to being able to make SET payments is to register the user's biometric with the netpage registration server. This is done in a controlled environment, for example a bank, where the biometric can be captured at the same time as the user's identity is verified. The biometric is captured and stored in the registration database, linked to the user's record. The user's photograph is also optionally captured and linked to the record. The SET cardholder registration process is completed, and the resulting private signature key and certificate are stored in the database. The user's payment card information is also stored, giving the netpage registration server enough information to act as the user's proxy in any SET payment transaction.

When the user eventually supplies the biometric to complete a payment, for example by signing a netpage order form, the printer securely transmits the order information, the pen ID and the biometric data to the netpage registration server. The server verifies the biometric with respect to the user identified by the pen ID, and from then on acts as the user's proxy in completing the SET payment transaction.

4.3 MICRO-PAYMENTS

The netpage system includes a mechanism for micro-payments, to allow the user to be conveniently charged for printing low-cost documents on demand and for copying copyright documents, and possibly also to allow the user to be reimbursed for expenses incurred in printing advertising material. The latter depends on the level of subsidy already provided to the user.

When the user registers for e-commerce, a network account is established which aggregates micro-payments. The user receives a statement on a regular basis, and can settle any outstanding debit balance using the standard payment mechanism.

The network account can be extended to aggregate subscription fees for periodicals, which would also otherwise be presented to the user in the form of individual statements.

4.4 TRANSACTIONS

When a user requests a netpage in a particular application context, the application is able to embed a user-specific transaction ID 55 in the page. Subsequent input through the page is tagged with the transaction ID, and the application is thereby able to establish an appropriate context for the user's input.

When input occurs through a page which is not user-specific, however, the application must use the user's unique identity to establish a context. A typical example involves adding items from a pre-printed catalog page to the user's virtual "shopping cart". To protect the user's privacy, however, the unique user ID 60 known to the netpage system is not divulged to applications. This is to prevent different application providers from easily correlating independently accumulated behavioral data.

The netpage registration server instead maintains an anonymous relationship between a user and an application via a unique alias ID 65, as shown in Figure 24. Whenever the user activates a hyperlink tagged with the "registered" attribute, the netpage page server asks the netpage registration server to translate the associated application ID 64, together with the pen ID 61, into an alias ID 65. The alias ID is then submitted to the hyperlink's application.

The application maintains state information indexed by alias ID, and is able to retrieve user-specific state information without knowledge of the global identity of the user.

The system also maintains an independent certificate and private signature key for each of a user's applications, to allow it to sign application transactions on behalf of the user using only application-specific information.

To assist the system in routing product bar code (e.g. UPC) and similar product-item-related "hyperlink" activations, the system records a favorite application on behalf of the user for any number of product types. For example, a user may nominate Amazon as their favorite bookseller, while a different user may nominate Barnes and Noble. When the

first user requests book-related information, e.g. via a printed book review or via an actual book, they are provided with the information by Amazon.

Each application is associated with an application provider, and the system maintains an account on behalf of each application provider, to allow it to credit and debit the provider for click-through fees etc.

An application provider can be a publisher of periodical subscribed content. The system records the user's willingness to receive the subscribed publication, as well as the expected frequency of publication.

5 COMMUNICATIONS PROTOCOLS

A communications protocol defines an ordered exchange of messages between entities. In the netpage system, entities such as pens, printers and servers utilise a set of defined protocols to cooperatively handle user interaction with the netpage system.

Each protocol is illustrated by way of a sequence diagram in which the horizontal dimension is used to represent message flow and the vertical dimension is used to represent time. Each entity is represented by a rectangle containing the name of the entity and a vertical column representing the lifeline of the entity. During the time an entity exists, the lifeline is shown as a dashed line. During the time an entity is active, the lifeline is shown as a double line. Because the protocols considered here do not create or destroy entities, lifelines are generally cut short as soon as an entity ceases to participate in a protocol.

5.1 SUBSCRIPTION DELIVERY PROTOCOL

A preferred embodiment of a subscription delivery protocol is shown in Figure 40.

A large number of users may subscribe to a periodical publication. Each user's edition may be laid out differently, but many users' editions will share common content such as text objects and image objects. The subscription delivery protocol therefore delivers document structures to individual printers via pointcast, but delivers shared content objects via multicast.

The application (i.e. publisher) first obtains a document ID 51 for each document from an ID server 12. It then sends each document structure, including its document ID and page descriptions, to the page server 10 responsible for the document's newly allocated ID. It includes its own application ID 64, the subscriber's alias ID 65, and the relevant set of multicast channel names. It signs the message using its private signature key.

The page server uses the application ID and alias ID to obtain from the registration server the corresponding user ID 60, the user's selected printer ID 62 (which may be explicitly selected for the application, or may be the user's default printer), and the application's certificate.

The application's certificate allows the page server to verify the message signature. The page server's request to the registration server fails if the application ID and alias ID don't together identify a subscription 808.

The page server then allocates document and page instance IDs and forwards the page descriptions, including page IDs 50, to the printer. It includes the relevant set of multicast channel names for the printer to listen to.

It then returns the newly allocated page IDs to the application for future reference.

Once the application has distributed all of the document structures to the subscribers' selected printers via the relevant page servers, it multicasts the various subsets of the shared objects on the previously selected multicast channels. Both page servers and printers monitor the appropriate multicast channels and receive their required content objects. They are then able to populate the previously pointcast document structures. This allows the page servers to add complete documents to their databases, and it allows the printers to print the documents.

5.2 HYPERLINK ACTIVATION PROTOCOL

A preferred embodiment of a hyperlink activation protocol is shown in Figure 42.

When a user clicks on a netpage with a netpage pen, the pen communicates the click to the nearest netpage printer 601. The click identifies the page and a location on the page. The printer already knows the ID 61 of the pen from the pen connection protocol.

The printer determines, via the DNS, the network address of the page server 10a handling the particular page ID 50. The address may already be in its cache if the user has recently interacted with the same page. The printer then forwards the pen ID, its own printer ID 62, the page ID and click location to the page server.

The page server loads the page description 5 identified by the page ID and determines which input element's zone 58, if any, the click lies in. Assuming the relevant input element is a hyperlink element 844, the page server then obtains the associated application ID 64 and link ID 54, and determines, via the DNS, the network address of the application server hosting the application 71.

The page server uses the pen ID 61 to obtain the corresponding user ID 60 from the registration server 11, and then allocates a globally unique hyperlink request ID 52 and builds a hyperlink request 934. The hyperlink request class diagram is shown in Figure 41. The hyperlink request records the IDs of the requesting user and printer, and identifies the clicked hyperlink instance 862. The page server then sends its own server ID 53, the hyperlink request ID, and the link ID to the application.

The application produces a response document according to application-specific logic, and obtains a document ID 51 from an ID server 12. It then sends the document to the page server 10b responsible for the document's newly allocated ID, together with the requesting page server's ID and the hyperlink request ID.

The second page server sends the hyperlink request ID and application ID to the first page server to obtain the corresponding user ID and printer ID 62. The first page server rejects the request if the hyperlink request has expired or is for a different application.

The second page server allocates document instance and page IDs 50, returns the newly allocated page IDs to the application, adds the complete document to its own database, and finally sends the page descriptions to the requesting printer.

The hyperlink instance may include a meaningful transaction ID 55, in which case the first page server includes the transaction ID in the message sent to the application. This allows the application to establish a transaction-specific context for the hyperlink activation.

If the hyperlink requires a user alias, i.e. its "alias required" attribute is set, then the first page server sends both the pen ID 61 and the hyperlink's application ID 64 to the registration server 11 to obtain not just the user ID corresponding to the pen ID but also the alias ID 65 corresponding to the application ID and the user ID. It includes the alias ID in the message sent to the application, allowing the application to establish a user-specific context for the hyperlink activation.

5.3 Handwriting Recognition Protocol

When a user draws a stroke on a netpage with a netpage pen, the pen communicates the stroke to the nearest netpage printer. The stroke identifies the page and a path on the page.

The printer forwards the pen ID 61, its own printer ID 62, the page ID 50 and stroke path to the page server 10 in the usual way.

The page server loads the page description 5 identified by the page ID and determines which input element's zone 58, if any, the stroke intersects. Assuming the relevant input element is a text field 878, the page server appends the stroke to the text field's digital ink.

After a period of inactivity in the zone of the text field, the page server sends the pen ID and the pending strokes to the registration server 11 for interpretation. The registration server identifies the user corresponding to the pen, and uses the user's accumulated handwriting model 822 to interpret the strokes as handwritten text. Once it has converted the strokes to text, the registration server returns the text to the requesting page server. The page server appends the text to the text value of the text field.

5.4 SIGNATURE VERIFICATION PROTOCOL

Assuming the input element whose zone the stroke intersects is a signature field 880, the page server 10 appends the stroke to the signature field's digital ink.

After a period of inactivity in the zone of the signature field, the page server sends the pen ID 61 and the pending strokes to the registration server 11 for verification. It also sends the application ID 64 associated with the form of which the signature field is part, as well as the form ID 56 and the current data content of the form. The registration server identifies the user corresponding to the pen, and uses the user's dynamic signature biometric 818 to verify the strokes as the user's signature. Once it has verified the signature, the registration server uses the application ID 64 and user ID 60 to identify the user's application-specific private signature key. It then uses the key to generate a digital signature of the form data, and returns the digital signature to the requesting page server. The page server assigns the digital signature to the signature field and sets the associated form's status to frozen.

The digital signature includes the alias ID 65 of the corresponding user. This allows a single form to capture multiple users' signatures.

5.5 FORM SUBMISSION PROTOCOL

A preferred embodiment of a form submission protocol is shown in Figure 43.

Form submission occurs via a form hyperlink activation. It thus follows the protocol defined in Section 5.2, with some form-specific additions.

In the case of a form hyperlink, the hyperlink activation message sent by the page server 10 to the application 71 also contains the form ID 56 and the current data content of

the form. If the form contains any signature fields, then the application verifies each one by extracting the alias ID 65 associated with the corresponding digital signature and obtaining the corresponding certificate from the registration server 11.

6 . . NETPAGE PEN DESCRIPTION

6.1 PEN MECHANICS

Referring to Figures 8 and 9, the pen, generally designated by reference numeral 101, includes a housing 102 in the form of a plastics moulding having walls 103 defining an interior space 104 for mounting the pen components. The pen top 105 is in operation rotatably mounted at one end 106 of the housing 102. A semi-transparent cover 107 is secured to the opposite end 108 of the housing 102. The cover 107 is also of moulded plastics, and is formed from semi-transparent material in order to enable the user to view the status of the LED mounted within the housing 102. The cover 107 includes a main part 109 which substantially surrounds the end 108 of the housing 102 and a projecting portion 110 which projects back from the main part 109 and fits within a corresponding slot 111 formed in the walls 103 of the housing 102. A radio antenna 112 is mounted behind the projecting portion 110, within the housing 102. Screw threads 113 surrounding an aperture 113A on the cover 107 are arranged to receive a metal end piece 114, including corresponding screw threads 115. The metal end piece 114 is removable to enable ink cartridge replacement.

Also mounted within the cover 107 is a tri-color status LED 116 on a flex PCB 117. The antenna 112 is also mounted on the flex PCB 117. The status LED 116 is mounted at the top of the pen 101 for good all-around visibility.

The pen can operate both as a normal marking ink pen and as a non-marking stylus. An ink pen cartridge 118 with nib 119 and a stylus 120 with stylus nib 121 are mounted side by side within the housing 102. Either the ink cartridge nib 119 or the stylus nib 121 can be brought forward through open end 122 of the metal end piece 114, by rotation of the pen top 105. Respective slider blocks 123 and 124 are mounted to the ink cartridge 118 and stylus 120, respectively. A rotatable cam barrel 125 is secured to the pen top 105 in operation and arranged to rotate therewith. The cam barrel 125 includes a cam 126 in the form of a slot within the walls 181 of the cam barrel. Cam followers 127 and 128 projecting from slider blocks 123 and 124 fit within the cam slot 126. On rotation of the cam barrel 125, the slider blocks 123 or 124 move relative to each other to project either the pen nib 119 or stylus nib 121 out through the hole 122 in the metal end piece 114. The pen 101 has three states of operation. By turning the top 105 through 90° steps, the three states are:

- stylus 120 nib 121 out

- ink cartridge 118 nib 119 out, and
- neither ink cartridge 118 nib 119 out nor stylus 120 nib 121 out

A second flex PCB 129, is mounted on an electronics chassis 130 which sits within the housing 102. The second flex PCB 129 mounts an infrared LED 131 for providing infrared radiation for projection onto the surface. An image sensor 132 is provided mounted on the second flex PCB 129 for receiving reflected radiation from the surface. The second flex PCB 129 also mounts a radio frequency chip 133, which includes an RF transmitter and RF receiver, and a controller chip 134 for controlling operation of the pen 101. An optics block 135 (formed from moulded clear plastics) sits within the cover 107 and projects an infrared beam onto the surface and receives images onto the image sensor 132. Power supply wires 136 connect the components on the second flex PCB 129 to battery contacts 137 which are mounted within the cam barrel 125. A terminal 138 connects to the battery contacts 137 and the cam barrel 125. A three volt rechargeable battery 139 sits within the cam barrel 125 in contact with the battery contacts. An induction charging coil 140 is mounted about the second flex PCB 129 to enable recharging of the battery 139 via induction. The second flex PCB 129 also mounts an infrared LED 143 and infrared photodiode 144 for detecting displacement in the cam barrel 125 when either the stylus 120 or the ink cartridge 118 is used for writing, in order to enable a determination of the force being applied to the surface by the pen nib 119 or stylus nib 121. The IR photodiode 144 detects light from the IR LED 143 via reflectors (not shown) mounted on the slider blocks 123 and 124.

Rubber grip pads 141 and 142 are provided towards the end 108 of the housing 102 to assist gripping the pen 101, and top 105 also includes a clip 142 for clipping the pen 101 to a pocket.

6.2 PEN CONTROLLER

The pen 101 is arranged to determine the position of its nib (stylus nib 121 or ink cartridge nib 119) by imaging, in the infrared spectrum, an area of the surface in the vicinity of the nib. It records the location data from the nearest location tag, and is arranged to calculate the distance of the nib 121 or 119 from the location tag utilising optics 135 and controller chip 134. The controller chip 134 calculates the orientation of the pen and the nib-to-tag distance from the perspective distortion observed on the imaged tag.

Utilising the RF chip 133 and antenna 112 the pen 101 can transmit the digital ink data (which is encrypted for security and packaged for efficient transmission) to the computing system.

When the pen is in range of a receiver, the digital ink data is transmitted as it is formed. When the pen 101 moves out of range, digital ink data is buffered within the pen 101 (the pen 101 circuitry includes a buffer arranged to store digital ink data for approximately 12 minutes of the pen motion on the surface) and can be transmitted later.

The controller chip 134 is mounted on the second flex PCB 129 in the pen 101. Figure 10 is a block diagram illustrating in more detail the architecture of the controller chip 134. Figure 10 also shows representations of the RF chip 133, the image sensor 132, the tri-color status LED 116, the IR illumination LED 131, the IR force sensor LED 143, and the force sensor photodiode 144.

The pen controller chip 134 includes a controlling processor 145. Bus 146 enables the exchange of data between components of the controller chip 134. Flash memory 147 and a 512 KB DRAM 148 are also included. An analog-to-digital converter 149 is arranged to convert the analog signal from the force sensor photodiode 144 to a digital signal.

An image sensor interface 152 interfaces with the image sensor 132. A transceiver controller 153 and base band circuit 154 are also included to interface with the RF chip 133 which includes an RF circuit 155 and RF resonators and inductors 156 connected to the antenna 112.

The controlling processor 145 captures and decodes location data from tags from the surface via the image sensor 132, monitors the force sensor photodiode 144, controls the LEDs 116, 131 and 143, and handles short-range radio communication via the radio transceiver 153. It is a medium-performance (~40MHz) general-purpose RISC processor.

The processor 145, digital transceiver components (transceiver controller 153 and baseband circuit 154), image sensor interface 152, flash memory 147 and 512KB DRAM 148 are integrated in a single controller ASIC. Analog RF components (RF circuit 155 and RF resonators and inductors 156) are provided in the separate RF chip.

The image sensor is a CCD or CMOS image sensor. Depending on tagging scheme, it has a size ranging from about 100x100 pixels to 200x200 pixels. Many miniature CMOS image sensors are commercially available, including the National Semiconductor LM9630.

The controller ASIC 134 enters a quiescent state after a period of inactivity when the pen 101 is not in contact with a surface. It incorporates a dedicated circuit 150 which monitors the force sensor photodiode 144 and wakes up the controller 134 via the power manager 151 on a pen-down event.

The radio transceiver communicates in the unlicensed 900MHz band normally used by cordless telephones, or alternatively in the unlicensed 2.4GHz industrial, scientific and medical (ISM) band, and uses frequency hopping and collision detection to provide interference-free communication.

In an alternative embodiment, the pen incorporates an Infrared Data Association (IrDA) interface for short-range communication with a base station or netpage printer.

In a further embodiment, the pen 101 includes a pair of orthogonal accelerometers mounted in the normal plane of the pen 101 axis. The accelerometers 190 are shown in Figures 9 and 10 in ghost outline.

The provision of the accelerometers enables this embodiment of the pen 101 to sense motion without reference to surface location tags, allowing the location tags to be sampled at a lower rate. Each location tag ID can then identify an object of interest rather than a position on the surface. For example, if the object is a user interface input element (e.g. a command button), then the tag ID of each location tag within the area of the input element can directly identify the input element.

The acceleration measured by the accelerometers in each of the x and y directions is integrated with respect to time to produce an instantaneous velocity and position.

Since the starting position of the stroke is not known, only relative positions within a stroke are calculated. Although position integration accumulates errors in the sensed acceleration, accelerometers typically have high resolution, and the time duration of a stroke, over which errors accumulate, is short.

7 NETPAGE PRINTER DESCRIPTION

7.1 PRINTER MECHANICS

The vertically-mounted netpage wallprinter 601 is shown fully assembled in Figure 11. It prints netpages on Letter/A4 sized media using duplexed 8½" Memjet™ print engines 602 and 603, as shown in Figures 12 and 12a. It uses a straight paper path with the paper 604 passing through the duplexed print engines 602 and 603 which print both sides of a sheet simultaneously, in full color and with full bleed.

An integral binding assembly 605 applies a strip of glue along one edge of each printed sheet, allowing it to adhere to the previous sheet when pressed against it. This creates a final bound document 618 which can range in thickness from one sheet to several hundred sheets.

The replaceable ink cartridge 627, shown in Figure 13 coupled with the duplexed print engines, has bladders or chambers for storing fixative, adhesive, and cyan, magenta, yellow, black and infrared inks. The cartridge also contains a micro air filter in a base molding. The micro air filter interfaces with an air pump 638 inside the printer via a hose 639. This provides filtered air to the printheads to prevent ingress of micro particles into the Memjet™ printheads 350 which might otherwise clog the printhead nozzles. By incorporating the air filter within the cartridge, the operational life of the filter is effectively linked to the life of the cartridge. The ink cartridge is a fully recyclable product with a capacity for printing and gluing 3000 pages (1500 sheets).

Referring to Figure 12, the motorized media pick-up roller assembly 626 pushes the top sheet directly from the media tray past a paper sensor on the first print engine 602 into the duplexed Memjet™ printhead assembly. The two Memjet™ print engines 602 and 603 are mounted in an opposing in-line sequential configuration along the straight paper path. The paper 604 is drawn into the first print engine 602 by integral, powered pick-up rollers 626. The position and size of the paper 604 is sensed and full bleed printing commences. Fixative is printed simultaneously to aid drying in the shortest possible time.

The paper exits the first Memjet™ print engine 602 through a set of powered exit spike wheels (aligned along the straight paper path), which act against a rubberized roller. These spike wheels contact the 'wet' printed surface and continue to feed the sheet 604 into the second Memjet™ print engine 603.

Referring to Figures 12 and 12a, the paper 604 passes from the duplexed print engines 602 and 603 into the binder assembly 605. The printed page passes between a powered spike wheel axle 670 with a fibrous support roller and another movable axle with spike wheels and a momentary action glue wheel. The movable axle/glue assembly 673 is mounted to a metal support bracket and it is transported forward to interface with the powered axle 670 via gears by action of a camshaft. A separate motor powers this camshaft.

The glue wheel assembly 673 consists of a partially hollow axle 679 with a rotating coupling for the glue supply hose 641 from the ink cartridge 627. This axle 679 connects to a glue wheel, which absorbs adhesive by capillary action through radial holes. A molded housing 682 surrounds the glue wheel, with an opening at the front. Pivoting side moldings and sprung outer doors are attached to the metal bracket and hinge out sideways when the rest of the assembly 673 is thrust forward. This action exposes the glue wheel through the front of the molded housing 682. Tension springs close the assembly and effectively cap the glue wheel during periods of inactivity.

As the sheet 604 passes into the glue wheel assembly 673, adhesive is applied to one vertical edge on the front side (apart from the first sheet of a document) as it is transported down into the binding assembly 605.

7.2 PRINTER CONTROLLER ARCHITECTURE

The netpage printer controller consists of a controlling processor 750, a factory-installed or field-installed network interface module 625, a radio transceiver (transceiver controller 753, baseband circuit 754, RF circuit 755, and RF resonators and inductors 756), dual raster image processor (RIP) DSPs 757, duplexed print engine controllers 760a and 760b, flash memory 658, and 64MB of DRAM 657, as illustrated in Figure 14.

The controlling processor handles communication with the network 19 and with local wireless netpage pens 101, senses the help button 617, controls the user interface LEDs 613-616, and feeds and synchronizes the RIP DSPs 757 and print engine controllers 760. It consists of a medium-performance general-purpose microprocessor. The controlling processor 750 communicates with the print engine controllers 760 via a high-speed serial bus 659.

The RIP DSPs rasterize and compress page descriptions to the netpage printer's compressed page format. Each print engine controller expands, dithers and prints page images to its associated Memjet™ printhead 350 in real time (i.e. at over 30 pages per minute). The duplexed print engine controllers print both sides of a sheet simultaneously.

The master print engine controller 760a controls the paper transport and monitors ink usage in conjunction with the master QA chip 665 and the ink cartridge QA chip 761.

The printer controller's flash memory 658 holds the software for both the processor 750 and the DSPs 757, as well as configuration data. This is copied to main memory 657 at boot time.

The processor 750, DSPs 757, and digital transceiver components (transceiver controller 753 and baseband circuit 754) are integrated in a single controller ASIC 656. Analog RF components (RF circuit 755 and RF resonators and inductors 756) are provided in a separate RF chip 762. The network interface module 625 is separate, since netpage printers allow the network connection to be factory-selected or field-selected. Flash memory 658 and the 2x256Mbit (64MB) DRAM 657 is also off-chip. The print engine controllers 760 are provided in separate ASICs.

A variety of network interface modules 625 are provided, each providing a netpage network interface 751 and optionally a local computer or network interface 752.

Netpage network Internet interfaces include POTS modems, Hybrid Fiber-Coax (HFC) cable modems, ISDN modems, DSL modems, satellite transceivers, current and next-generation cellular telephone transceivers, and wireless local loop (WLL) transceivers. Local interfaces include IEEE 1284 (parallel port), 10Base-T and 100Base-T Ethernet, USB and USB 2.0, IEEE 1394 (Firewire), and various emerging home networking interfaces. If an Internet connection is available on the local network, then the local network interface can be used as the netpage network interface.

The radio transceiver 753 communicates in the unlicensed 900MHz band normally used by cordless telephones, or alternatively in the unlicensed 2.4GHz industrial, scientific and medical (ISM) band, and uses frequency hopping and collision detection to provide interference-free communication.

The printer controller optionally incorporates an Infrared Data Association (IrDA) interface for receiving data "squirted" from devices such as netpage cameras. In an alternative embodiment, the printer uses the IrDA interface for short-range communication with suitably configured netpage pens.

7.2.1 Rasterization and Printing

Once the main processor 750 has received and verified the document's page layouts and page objects, it runs the appropriate RIP software on the DSPs 757.

The DSPs 757 rasterize each page description and compress the rasterized page image. The main processor stores each compressed page image in memory. The simplest way to load-balance multiple DSPs is to let each DSP rasterize a separate page. The DSPs can always be kept busy since an arbitrary number of rasterized pages can, in general, be stored in memory. This strategy only leads to potentially poor DSP utilization when rasterizing short documents.

Watermark regions in the page description are rasterized to a contone-resolution bi-level bitmap which is losslessly compressed to negligible size and which forms part of the compressed page image.

The infrared (IR) layer of the printed page contains coded netpage tags at a density of about six per inch. Each tag encodes the page ID, tag ID, and control bits, and the data content of each tag is generated during rasterization and stored in the compressed page image.

The main processor 750 passes back-to-back page images to the duplexed print engine controllers 760. Each print engine controller 760 stores the compressed page image in its local memory, and starts the page expansion and printing pipeline. Page expansion and

printing is pipelined because it is impractical to store an entire 114MB bi-level CMYK+IR page image in memory.

7.2.2 Print Engine Controller

The page expansion and printing pipeline of the print engine controller 760 consists of a high speed IEEE 1394 serial interface 659, a standard JPEG decoder 763, a standard Group 4 Fax decoder 764, a custom halftoner/compositor unit 765, a custom tag encoder 766, a line loader/formatter unit 767, and a custom interface 768 to the Memjet™ printhead 350.

The print engine controller 360 operates in a double buffered manner. While one page is loaded into DRAM 769 via the high speed serial interface 659, the previously loaded page is read from DRAM 769 and passed through the print engine controller pipeline. Once the page has finished printing, the page just loaded is printed while another page is loaded.

The first stage of the pipeline expands (at 763) the JPEG-compressed contone CMYK layer, expands (at 764) the Group 4 Fax-compressed bi-level black layer, and renders (at 766) the bi-level netpage tag layer according to the tag format defined in section 1.2, all in parallel. The second stage dithers (at 765) the contone CMYK layer and composites (at 765) the bi-level black layer over the resulting bi-level CMYK layer. The resultant bi-level CMYK+IR dot data is buffered and formatted (at 767) for printing on the Memjet™ printhead 350 via a set of line buffers. Most of these line buffers are stored in the off-chip DRAM. The final stage prints the six channels of bi-level dot data (including fixative) to the Memjet™ printhead 350 via the printhead interface 768.

When several print engine controllers 760 are used in unison, such as in a duplexed configuration, they are synchronized via a shared line sync signal 770. Only one print engine 760, selected via the external master/slave pin 771, generates the line sync signal 770 onto the shared line.

The print engine controller 760 contains a low-speed processor 772 for synchronizing the page expansion and rendering pipeline, configuring the printhead 350 via a low-speed serial bus 773, and controlling the stepper motors 675, 676.

In the 8½" versions of the netpage printer, the two print engines each prints 30 Letter pages per minute along the long dimension of the page (11"), giving a line rate of 8.8 kHz at 1600 dpi. In the 12" versions of the netpage printer, the two print engines each prints 45 Letter pages per minute along the short dimension of the page (8½"), giving a line rate of

10.2 kHz. These line rates are well within the operating frequency of the Memjet™ print-head, which in the current design exceeds 30 kHz.

8 PRODUCT TAGGING

Automatic identification refers to the use of technologies such as bar codes, magnetic stripe cards, smartcards, and RF transponders, to (semi-)automatically identify objects to data processing systems without manual keying.

For the purposes of automatic identification, a product item is commonly identified by a 12-digit Universal Product Code (UPC), encoded machine-readably in the form of a printed bar code. The most common UPC numbering system incorporates a 5-digit manufacturer ID and a 5-digit item number. Because of its limited precision, a UPC is used to identify a class of product rather than an individual product item. The Uniform Code Council and EAN International define and administer the UPC and related codes as subsets of the 14-digit Global Trade Item Number (GTIN).

Within supply chain management, there is considerable interest in expanding or replacing the UPC scheme to allow individual product items to be uniquely identified and thereby tracked. Individual item tagging can reduce "shrinkage" due to lost, stolen or spoiled goods, improve the efficiency of demand-driven manufacturing and supply, facilitate the profiling of product usage, and improve the customer experience.

There are two main contenders for individual item tagging: optical tags in the form of so-called two-dimensional bar codes, and radio frequency identification (RFID) tags. For a detailed description of RFID tags, refer to Klaus Finkenzeller, *RFID Handbook*, John Wiley & Son (1999), the contents of which are herein incorporated by cross-reference. Optical tags have the advantage of being inexpensive, but require optical line-of-sight for reading. RFID tags have the advantage of supporting omnidirectional reading, but are comparatively expensive. The presence of metal or liquid can seriously interfere with RFID tag performance, undermining the omnidirectional reading advantage. Passive (reader-powered) RFID tags are projected to be priced at 10 cents each in multi-million quantities by the end of 2003, and at 5 cents each soon thereafter, but this still falls short of the sub-one-cent industry target for low-price items such as grocery. The read-only nature of most optical tags has also been cited as a disadvantage, since status changes cannot be written to a tag as an item progresses through the supply chain. However, this disadvantage is mitigated by the fact that a read-only tag can refer to information maintained dynamically on a network.

The Massachusetts Institute of Technology (MIT) Auto-ID Center has developed a standard for a 96-bit Electronic Product Code (EPC), coupled with an Internet-based Object

Naming Service (ONS) and a Product Markup Language (PML). Once an EPC is scanned or otherwise obtained, it is used to look up, possibly via the ONS, matching product information portably encoded in PML. The EPC consists of an 8-bit header, a 28-bit EPC manager, a 24-bit object class, and a 36-bit serial number. For a detailed description of the EPC, refer to Brock, D.L., *The Electronic Product Code (EPC)*, MIT Auto-ID Center (January 2001), the contents of which are herein incorporated by cross-reference. The Auto-ID Center has defined a mapping of the GTIN onto the EPC to demonstrate compatibility between the EPC and current practices Brock, D.L., *Integrating the Electronic Product Code (EPC) and the Global Trade Item Number (GTIN)*, MIT Auto-ID Center (November 2001), the contents of which are herein incorporated by cross-reference.

Although EPCs can be encoded and carried in many forms, the Auto-ID Center strongly advocates the use of low-cost passive RFID tags to carry EPCs, and has defined a 64-bit version of the EPC to allow the cost of RFID tags to be minimized in the short term. For detailed description of low-cost RFID tag characteristics, refer to Sarma, S., *Towards the 5c Tag*, MIT Auto-ID Center (November 2001), the contents of which are herein incorporated by cross-reference. For a description of a commercially-available low-cost passive RFID tag, refer to *915 MHz RFID Tag*, Alien Technology (2002), the contents of which are herein incorporated by cross-reference. For detailed description of the 64-bit EPC, refer to Brock, D.L., *The Compact Electronic Product Code*, MIT Auto-ID Center (November 2001), the contents of which are herein incorporated by cross-reference.

EPCs are intended not just for unique item-level tagging and tracking, but also for case-level and pallet-level tagging, and for tagging of other logistic units of shipping and transportation such as containers and trucks. The distributed PML database records dynamic relationships between items and higher-level containers in the packaging, shipping and transportation hierarchy.

8.1 OMNITAGGING IN THE SUPPLY CHAIN

Using an invisible (e.g. infrared) tagging scheme to uniquely identify a product item has the significant advantage that it allows the entire surface of a product to be tagged, or a significant portion thereof, without impinging on the graphic design of the product's packaging or labelling. If the entire product surface is tagged, then the orientation of the product doesn't affect its ability to be scanned, i.e. a significant part of the line-of-sight disadvantage of a visible bar code is eliminated. Furthermore, since the tags are small and massively replicated, label damage no longer prevents scanning.

Omnitagging, then, consists of covering a large proportion of the surface of a product item with optically-readable invisible tags. Each omnitag uniquely identifies the product item on which it appears. The omnitag may directly encode the product code (e.g. EPC) of the item, or may encode a surrogate ID which in turn identifies the product code via a database lookup. Each omnitag also optionally identifies its own position on the surface of the product item, to provide the downstream consumer benefits of netpage interactivity described earlier.

Omnitags are applied during product manufacture and/or packaging using digital printers. These may be add-on infrared printers which print the omnitags after the text and graphics have been printed by other means, or integrated color and infrared printers which print the omnitags, text and graphics simultaneously. Digitally-printed text and graphics may include everything on the label or packaging, or may consist only of the variable portions, with other portions still printed by other means.

8.2 OMNITAGGING

As shown in Figure 18, a product's unique item ID 215 may be seen as a special kind of unique object ID 210. The Electronic Product Code (EPC) 220 is one emerging standard for an item ID. An item ID typically consists of a product ID 214 and a serial number 213. The product ID identifies a class of product, while the serial number identifies a particular instance of that class, i.e. an individual product item. The product ID in turn typically consists of a manufacturer ID 211 and a product class number 212. The best-known product ID is the EAN.UCC Universal Product Code (UPC) 221 and its variants.

As shown in Figure 19, an omnitag 202 encodes a page ID (or region ID) 50 and a two-dimensional (2D) position 86. The region ID identifies the surface region containing the tag, and the position identifies the tag's position within the two-dimensional region. Since the surface in question is the surface of a physical product item 201, it is useful to define a one-to-one mapping between the region ID and the unique object ID 210, and more specifically the item ID 215, of the product item. Note, however, that the mapping can be many-to-one without compromising the utility of the omnitag. For example, each panel of a product item's packaging could have a different region ID 50. Conversely, the omnitag may directly encode the item ID, in which case the region ID contains the item ID, suitably prefixed to decouple item ID allocation from general netpage region ID allocation. Note that the region ID uniquely distinguishes the corresponding surface region from all other surface regions identified within the global netpage system.

The item ID 215 is preferably the EPC 220 proposed by the Auto-ID Center, since this provides direct compatibility between omnitags and EPC-carrying RFID tags.

In Figure 19 the position 86 is shown as optional. This is to indicate that much of the utility of the omnitag in the supply chain derives from the region ID 50, and the position may be omitted if not desired for a particular product.

For interoperability with the netpage system, an omnitag 202 is a netpage tag 4, i.e. it has the logical structure, physical layout and semantics of a netpage tag.

When a netpage sensing device such as the netpage pen 101 images and decodes an omnitag, it uses the position encoded in the tag, and the position and orientation of the tag in its field of view, to compute its own position relative to the tag and hence relative to the region containing the tag. As the sensing device is moved relative to an omnitagged surface region, it is thereby able to track its own motion relative to the region and generate a set of timestamped position samples representative of its time-varying path. When the sensing device is a pen, then the path consists of a sequence of strokes, with each stroke starting when the pen makes contact with the surface, and ending when the pen breaks contact with the surface.

When a stroke is forwarded to the page server 10 responsible for the region ID, the server retrieves a description of the region keyed by region ID, and interprets the stroke in relation to the description. For example, if the description includes a hyperlink and the stroke intersects the zone of the hyperlink, then the server may interpret the stroke as a designation of the hyperlink and activate the hyperlink.

8.2.1 Item Id Management

As previously described, a structured item ID typically has a three-level encoding, consisting of a manufacturer ID, a product class number, and a serial number. In the EPC the manufacturer ID corresponds to the manager ID. Manufacturer ids are assigned to particular manufacturers 235 by a governing body such the Uniform Code Council (UCC). Within the scope of each manufacturer ID the manufacturer 235 assigns product class numbers to particular product classes 236, and within the scope of each product class number the manufacturer assigns serial numbers to individual product items 237. Each assignor in the assignment hierarchy ensures that each component of the item ID is assigned uniquely, with the end result that an item ID uniquely identifies a single product item. Each assigned item ID component is robustly recorded to ensure unique assignment, and subsequently becomes a database key to details about the corresponding manufacturer, product or item. At the product level this information may include the product's description, dimensions,

weight and price, while at the item level it may include the item's expiry date and place of manufacture.

As shown in Figure 20, a collection of related product classes may be recorded as a single product type 238, identified by a unique product type ID 217. This provides the basis for mapping a scanned or otherwise obtained product ID 214 (or the product ID portion of a scanned or otherwise obtained item ID 215) to a product type 238. This in turn allows a favorite application 828 for that product type to be identified for a particular netpage user 800, as shown in Figure 24.

As a product item moves through the supply chain, status information is ideally maintained in a globally accessible database, keyed by the item ID. This information may include the item's dynamic position in the packaging, shipping and transportation hierarchy, its location on a store shelf, and ultimately the date and time of its sale and the recipient of that sale. In a packaging, shipping and transportation hierarchy, higher level units such as cases, pallets, shipping containers and trucks all have their own item ids, and this provides the basis for recording the dynamic hierarchy in which the end product item participates. Note that the concept of an item also extends to a sub-component of an assembly or a component or element of a saleable product.

Figure 20 shows the product description hierarchy corresponding to the structure of the item id; the product item's dynamic participation in a dynamic packaging, shipping and transportation hierarchy; and the product item's dynamic ownership. As the figure shows, a container 231 (e.g. case, pallet, shipping container, or truck) is a special case of an uniquely identified object 230. The fact that the container is holding, or has held, a particular object for the duration of some time interval is represented by the time-stamped object location, wherein the end time remains unspecified until the container ceases to hold the item. The object-container relationship is recursive, allowing it to represent an arbitrary dynamic hierarchy. Clearly this representation can be expanded to record the time-varying relative or absolute geographic location of an object.

The fact that an entity 232 owns, or has owned, a particular object for the duration of some time interval is represented by the time-stamped object ownership 233, wherein the end time remains unspecified until the entity ceases to own the item. The owning entity 232 may represent a netpage user 800, e.g. when a netpage user purchases a product item and the sale is recorded.

As shown in Figure 56, a physical product item 201 is recorded as a product item 237 by a product server 251. A product item may be recorded in multiple product servers, managed by different participants in the supply chain such as manufacturers, distributors

and retailers. However, benefits accrue from providing a unified view of a product item, even if the unified view is provided virtually.

To foster interoperability between different supply chain participants and between disparate systems which may want to query and update both static and dynamic item information, such information interchanges are ideally performed using a standard representation. The MIT Auto-ID Center's Physical Markup Language (PML) is an example of a standard representation designed for this purpose. For a detailed description of PML, refer to Brock, D.L. *et al.*, *The Physical Markup Language*, MIT Auto-ID Center (June 2001), the contents of which are herein incorporated by cross-reference.

8.2.2 Region Id Management

An unstructured ID such as the region ID 50 may be assigned on demand through a multi-level assignment hierarchy with a single root node. Lower-level assignors obtain blocks of ids from higher-level assignors on demand. Unlike with structured ID assignment, these blocks correspond to arbitrary ranges (or even sets) of ids, rather than to ids with fixed prefixes. Again, each assignor in the assignment hierarchy ensures that blocks of ids and individual ids are assigned uniquely. The region ID subsequently becomes a database key to information about the region. In the netpage system, this information includes a full description of the graphical and interactive elements which appear in the region. Graphical elements may include such things as text flows, text and images. Interactive elements may include such things as buttons, hyperlinks, checkboxes, drawing fields, text fields and signature fields.

8.3 OMNITAG PRINTING

An omnitag printer is a digital printer which prints omnitags onto the label, packaging or actual surface of a product before, during or after product manufacture and/or assembly. It is a special case of a netpage printer 601. It is capable of printing a continuous pattern of omnitags onto a surface, typically using a near-infrared-absorptive ink. In high-speed environments, the printer includes hardware which accelerates tag rendering. This typically includes real-time Reed-Solomon encoding of variable tag data such as tag position, and real-time template-based rendering of the actual tag pattern at the dot resolution of the printhead.

The printer may be an add-on infrared printer which prints the omnitags after text and graphics have been printed by other means, or an integrated color and infrared printer which prints the omnitags, text and graphics simultaneously. Digitally-printed text and

graphics may include everything on the label or packaging, or may consist only of the variable portions, with other portions still printed by other means. Thus an omnitag printer with an infrared and black printing capability can displace an existing digital printer used for variable data printing, such as a conventional thermal transfer or inkjet printer.

For the purposes of the following discussion, any reference to printing onto an item label is intended to include printing onto the item packaging in general, or directly onto the item surface. Furthermore, any reference to an item ID 215 is intended to include a region ID 50 (or collection of per-panel region ids), or a component thereof.

The printer is typically controlled by a host computer, which supplies the printer with fixed and/or variable text and graphics as well as item ids for inclusion in the omnitags. The host may provide real-time control over the printer, whereby it provides the printer with data in real time as printing proceeds. As an optimisation, the host may provide the printer with fixed data before printing begins, and only provide variable data in real time. The printer may also be capable of generating per-item variable data based on parameters provided by the host. For example, the host may provide the printer with a base item ID prior to printing, and the printer may simply increment the base item ID to generate successive item ids. Alternatively, memory in the ink cartridge or other storage medium inserted into the printer may provide a source of unique item ids, in which case the printer reports the assignment of items ids to the host computer for recording by the host.

Alternatively still, the printer may be capable of reading a pre-existing item ID from the label onto which the omnitags are being printed, assuming the unique ID has been applied in some form to the label during a previous manufacturing step. For example, the item ID may already be present in the form of a visible 2D bar code, or encoded in an RFID tag. In the former case the printer can include an optical bar code scanner. In the latter case it can include an RFID reader.

The printer may also be capable of rendering the item ID in other forms. For example, it may be capable of printing the item ID in the form of a 2D bar code, or of printing the product ID component of the item ID in the form of a 1D bar code, or of writing the item ID to a writable or write-once RFID tag.

8.4 OMNITAG SCANNING

Item information typically flows to the product server in response to situated scan events, e.g. when an item is scanned into inventory on delivery; when the item is placed on a retail shelf; and when the item is scanned at point of sale. Both fixed and hand-held scanners may be used to scan omnitagged product items, using both laser-based 2D scanning

and 2D image-sensor-based scanning, using similar or the same techniques as employed in the netpage pen.

As shown in Figure 57, both a fixed scanner 254 and a hand-held scanner 252 communicate scan data to the product server 251. The product server may in turn communicate product item event data to a peer product server (not shown), or to a product application server 250, which may implement sharing of data with related product servers. For example, stock movements within a retail store may be recorded locally on the retail store's product server, but the manufacturer's product server may be notified once a product item is sold.

8.5 OMNITAG-BASED NETPAGE INTERACTIONS

A product item whose labelling, packaging or actual surface has been omnitagged provides the same level of interactivity as any other netpage.

There is a strong case to be made for netpage-compatible product tagging. Netpage turns any printed surface into a finely differentiated graphical user interface akin to a Web page, and there are many applications which map nicely onto the surface of a product. These applications include obtaining product information of various kinds (nutritional information; cooking instructions; recipes; related products; use-by dates; servicing instructions; recall notices); playing games; entering competitions; managing ownership (registration; query, such as in the case of stolen goods; transfer); providing product feedback; messaging; and indirect device control. If, on the other hand, the product tagging is undifferentiated, such as in the case of an undifferentiated 2D barcode or RFID-carried item ID, then the burden of information navigation is transferred to the information delivery device, which may significantly increase the complexity of the user experience or the required sophistication of the delivery device user interface.

8.5.1 Product Registration

An omnitagged product can contain a <register> button which, when activated with a netpage pen, registers the netpage user as the owner of the product. The user's contact information, which is already recorded on the netpage system, can be automatically transmitted to the product manufacturer who can record it in their customer database. The registration process can automatically add the manufacturer to the user's e-mail contact list, thus allowing the manufacturer to send the user e-mail relevant to the product, such as related special offers, recall notices, etc. If the manufacturer abuses their e-mail privileges, the user can bar them in the usual way.

8.5.2 Product Information via Product ID

Some of the benefits of omnitagging products can be gained by enhancing the netpage pen to decode UPC bar codes. Alternatively a UPC bar code scanner can netpage-enabled. When the netpage system receives a scanned UPC, it forwards a request to a default or favorite application for that product type (as described earlier), and this in turn elicits product information from the application, such as in the form of a printed netpage. The product page can also include the facility to enter the serial number of the product item and register the user's ownership of it via a <register> button. Product manufacturers can thus gain the benefits of netpage linking for their entire installed base of products without making alterations to the products themselves.

8.5.3 Context-Specific Product Help

If the entire surface of a product is omnitagged, then pressing on any part of the surface with a netpage pen can then elicit product-specific help. The help is either specific to the area pressed, or relates to the product as a whole. Thus the user of the product has instant access to helpful information about specific features of a product as well as the product as a whole. Each feature-specific help page can be linked to the entire product manual.

8.5.4 Product Ownership Tracking

If the entire surface of a product is omnitagged, then pressing on any part of the surface with a netpage pen can elicit a description of the product and its current ownership. After the product is purchased, pressing on any part of the surface can automatically register the product in the name of the owner of the netpage pen. Anyone can determine the ownership of a product offered for sale simply by pressing on any part of its surface with a Netpage Pen. Ownership may only be registered by a new owner if the current owner has relinquished ownership by signing the "sell" portion of the product's status page. This places the product in an "un-owned" state.

Product information and ownership is maintained either by the product manufacturer, as a service to its customers, or by a profit-oriented third party.

The shipping computer system of a product manufacturer can automatically transfer ownership of products from the manufacturer to the distributor or retailer, and so on down through the supply chain. The retail computer system of the retailer can automatically mark each sold item as free, or transfer ownership directly to the holder of the payment card used to pay for the product. The customer can also use a netpage pen at the point of sale to register immediate ownership of the product.

Traditional clearing-houses for stolen goods, such as pawn shops, can be required by law to check the ownership of all products presented to them. Since an omnitagged product has an invisible encoding on most or all of its surface, it is difficult for a thief to remove it or even tell if it has been successfully removed. Conversely, it is incumbent on a potential buyer of a product to ensure that a clean reading can be obtained from its surface so that its ownership can be indisputably established.

Where a product is leased or otherwise subject to complex or multiple ownership, the product registration database can reflect this and thus alert a potential buyer.

CONCLUSION

Although the invention has been described with reference to a number of specific examples, it will be appreciated by those skilled in the art that the invention can be embodied in many other forms.

CLAIMS

1. A method of facilitating an interaction between a user and a product item, the product item having an identity and the method including the steps of:

providing the user with an interface surface associated with the product item and containing information relating to the product item, the interface surface having disposed thereon coded data indicative of the identity of the product item and of a plurality of reference points of the interface surface;

receiving, in a computer system, indicating data from a sensing device regarding the identity of the product item and a position of the sensing device relative to the interface surface, the sensing device, when placed in an operative position relative to the interface surface, sensing at least some of the coded data in the vicinity of the sensing device and generating the indicating data using at least some of the sensed coded data; and

facilitating, in the computer system and with reference to the indicating data, the interaction between the user and the product item.

2. The method of claim 1 in which the interaction is associated with at least one zone of the interface surface and in which the method includes identifying, in the computer system and from the zone relative to which the sensing device is located, the interaction.

3. The method of claim 2 which includes:

receiving, in the computer system, data regarding movement of the sensing device relative to the interface surface, the sensing device sensing its movement relative to the interface surface using at least some of the coded data; and

identifying, in the computer system and from the movement being at least partially within the at least one zone, the interaction.

4. A method of facilitating an interaction between a user and a product item, the product item having an identity and the method including the steps of:

providing the user with an interface surface associated with the product item, the interface surface containing information relating to the product item and having disposed thereon coded data indicative of the identity of the product item and a parameter of the interaction;

receiving, in a computer system and from a sensing device, indicating data regarding the identity of both the product item and the parameter of the interaction, and movement data regarding movement of the sensing device relative to the interface surface, the sensing device, when moved relative to the interface surface, sensing at least some of the coded data

and generating the indicating data and the movement data using at least some of the sensed coded data; and

interpreting, in the computer system, the indicating data and the movement data as it relates to the interaction.

5. A method of facilitating an interaction between a user and a product item, the product item having an identity and the method including the steps of:

providing the user with an interface surface associated with the product item, the interface surface containing information relating to the product item and having disposed thereon coded data indicative of the identity of the product item;

receiving, in a computer system, user data from a sensing device regarding an identity of the user and indicating data regarding the identity of the product item, the sensing device containing the user data and, when placed in an operative position relative to the interface surface, sensing at least some of the coded data in the vicinity of the sensing device and generating the indicating data using at least some of the sensed coded data; and

facilitating, in the computer system and with reference to the user data and the indicating data, the interaction between the user and the product item.

6. The method of claim 5 in which the coded data is also indicative of a parameter of the interaction, and the method includes receiving, in the computer system, indicating data from the sensing device regarding the parameter of the interaction, the sensing device generating the indicating data using at least some of the sensed coded data.

7. The method of claim 6 which includes receiving, in the computer system, data from the sensing device regarding movement of the sensing device relative to the interface surface, the sensing device generating data regarding its own movement relative to the interface surface.

8. The method of any one of claims 1, 4 or 5 in which the interaction is selected from the group comprising:

- (a) providing product information about the product item to the user;
- (b) recording a purchase transaction relating to the product item;
- (c) recording a potential purchase transaction relating to the product item;
- (d) providing comparison information to the user, the comparison information comparing the product information about the product item with product information about another product item;
- (e) playing a game associated with the product item; and

(f) conducting a competition in relation to the product item.

9. A method according to claim 8 wherein the product information comprises information relating to any one of the product item's:

- (a) cost;
- (b) contents;
- (c) weight;
- (d) place of origin;
- (e) manufacturer;
- (f) date of manufacture;
- (g) date of packaging;
- (h) use-by date;
- (i) current owner; and
- (j) dimensions.

10. A method according to any one of claims 1, 4 or 5 wherein the interface surface is selected from the group comprising:

- (a) a label;
- (b) a package; and
- (c) a surface of the product item itself.

11. A method according to any one of claims 1, 4 or 5 wherein the coded data is substantially invisible to the average unaided human.

12. A method according to any one of claims 1, 4 or 5 wherein the coded data is disposed over a substantial portion of the interface surface.

13. A method according to claim 12 wherein the coded data is disposed over more than 20% of the interface surface.

14. A method according to claim 12 wherein the coded data is disposed over more than 90% of the interface surface.

15. The method of any one of claims 1, 4 or 5 wherein the method further comprises identifying, in the computer system, that the user has entered handwritten text data by means of the sensing device and effecting, in the computer system, an operation associated with the handwritten text data.

16. The method of claim 15 which includes converting, in the computer system, the handwritten text data to computer text.
17. The method of any one of claims 1, 4 or 5 wherein the method further comprises identifying, in the computer system, that the user has entered a handwritten signature by means of the sensing device and effecting, in the computer system, an operation associated with the handwritten signature.
18. The method of claim 17 which includes verifying, in the computer system, that the signature is that of the user.
19. The method of claim 17 in which the operation associated with the handwritten signature is associated with payment authorization.
20. The method of any one of claims 1, 4 or 5 wherein the method further comprises identifying, in the computer system, that the user has entered a hand-drawn picture by means of the sensing device and effecting, in the computer system, an operation associated with the hand-drawn picture.
21. The method of any one of claims 1, 4 or 5 in which:
a portion of the coded data is superimposed with a visual graphic, the visual graphic being at least partially indicative, to the user, of the interaction.
22. The method of claim 1 or claim 4 in which the sensing device contains an identifier which imparts a unique identity to the sensing device and identifies it as belonging to the user and in which the method includes monitoring, in the computer system, the identifier.
23. The method of any one of claims 1, 4 or 5 which includes providing all required information relating to the interaction in the interface surface to eliminate the need for a separate display device.
24. A system for facilitating an interaction between a user and a product item, the product item having an identity and the system including:
a interface surface associated with the product item and containing information relating to the product item, the interface surface having disposed thereon coded data indicative of the identity of the product item and of a plurality of reference points of the interface surface; and

a computer system adapted to facilitate the interaction in response to receiving indicating data from a sensing device, the indicating data being indicative of the identity of the product item and of a position of the sensing device relative to the interface surface, the sensing device, when placed in an operative position relative to the interface surface, sensing at least some of the coded data in the vicinity of the sensing device and generating the indicating data using at least some of the sensed coded data.

25. The system of claim 24 in which the interaction is associated with at least one zone of the interface surface.

26. The system of claim 24 which includes the sensing device, the sensing device sensing its movement relative to the interface surface using at least some of the coded data.

27. A system for facilitating an interaction between a user and a product item, the product item having an identity and the system including

an interface surface associated with the product item, the interface surface containing information relating to the product item and having disposed thereon coded data indicative of the identity of the product item and of a parameter of the interaction; and

a computer system for receiving, from a sensing device, indicating data regarding the identity of both the product item and the parameter of the interaction, and movement data regarding movement of the sensing device relative to the interface surface, and for interpreting the movement of the sensing device as it relates to the interaction, the sensing device, when moved relative to the interface surface, sensing at least some of the coded data and generating the indicating data and the movement data using at least some of the sensed coded data.

28. A system for facilitating an interaction between a user and a product item, the product item having an identity and the system including:

a interface surface associated with the product item, the interface surface containing information relating to the product item and having disposed thereon coded data indicative of the identity of the product item; and

a computer system adapted to receive from a sensing device user data regarding an identity of the user and indicating data regarding the identity of the product item, and for facilitating, with reference to the user data and the indicating data, the interaction between the user and the product item, the sensing device containing the user data and, when placed in an operative position relative to the interface surface, sensing at least some of the coded data in the vicinity of the sensing device and generating the indicating data using at least

some of the sensed coded data.

29. The system of claim 28 in which the coded data is also indicative of a parameter of the interaction, the computer system receiving indicating data from the sensing device regarding the parameter of the interaction, and the sensing device generating the indicating data using at least some of the coded data.

30. The system of claim 28 which includes the sensing device, the sensing device sensing its movement relative to the interface surface.

31. The system of any one of claims 24, 27 or 28 in which the interaction is selected from the group comprising:

- (a) providing product information about the product item to the user;
- (b) recording a purchase transaction relating to the product item;
- (c) recording a potential purchase transaction relating to the product item;
- (d) providing comparison information to the user, the comparison information comparing the product information about the product item with product information about another product item;
- (e) playing a game associated with the product item; and
- (f) conducting a competition in relation to the product item.

32. The system of claim 31 wherein the product information comprises information relating to any one of the product item's:

- (a) cost;
- (b) contents;
- (c) weight;
- (d) place of origin;
- (e) manufacturer;
- (f) date of manufacture;
- (g) date of packaging;
- (h) use-by date;
- (i) current owner; and
- (j) dimensions.

33. The system of any one of claims 24, 27 or 28 wherein the interface surface is selected from the group comprising:

- (a) a label;

- (b) a package; and
- (c) a surface of the product item itself.

34. The system of any one of claims 24, 27 or 28 wherein the coded data is substantially invisible to the average unaided human.

35. The system of any one of claims 24, 27 or 28 wherein the coded data is disposed over a substantial portion of the interface surface.

36. The system of claim 35 wherein the coded data is disposed over more than 20% of the interface surface.

37. The system of claim 35 wherein the coded data is disposed over more than 90% of the interface surface.

38. The system of any one of claims 24, 27 and 28 in which the sensing device includes a marking nib.

39. The system of any one of claims 24, 27 and 28 in which the sensing device contains an identifier which imparts a unique identity to the sensing device and identifies it as belonging to the user.

40. The method of any one of claims 1, 4 or 5 wherein the coded data is disposed on the interface surface in accordance with a layout, the layout having at least order n rotational symmetry, where n is at least two, the layout encoding an orientation codeword comprising a sequence of an integer multiple m of n symbols, where m is one or more, each encoded symbol being distributed at n locations about a center of rotational symmetry of the layout such that decoding the symbols at each of the n orientations of the layout produces n representations of the orientation codeword, each representation comprising a different cyclic shift of the orientation codeword and being indicative of the degree of rotation of the layout, and wherein the orientation codeword is fault tolerant.

41. The method of any one of claims 1, 4 or 5 wherein the coded data is disposed on the interface surface in accordance with a layout, the layout having at least order n rotational symmetry, where n is at least two, the layout including n identical sub-layouts rotated $1/n$ revolutions apart about a center of rotational symmetry of the layout, the coded data disposed in accordance with each sub-layout including rotation-indicating data that distin-

guishes the rotation of that sub-layout from the rotation of at least one other sub-layout within the layout.

42. The method of any one of claims 1, 4 or 5 wherein the coded data is disposed on the interface surface in accordance with a layout having n -fold rotational symmetry, where n is at least two, the layout including n identical first sub-layouts rotated $1/n$ revolutions apart about a center of rotational symmetry of the layout, the coded data disposed in accordance with each first sub-layout including rotation-indicating data that distinguishes the rotation of that first sub-layout from the rotation of at least one other first sub-layout within the layout.

43. The system of any one of claims 24, 27 or 28 wherein the coded data is disposed on the interface surface in accordance with a layout, the layout having at least order n rotational symmetry, where n is at least two, the layout encoding an orientation codeword comprising a sequence of an integer multiple m of n symbols, where m is one or more, each encoded symbol being distributed at n locations about a center of rotational symmetry of the layout such that decoding the symbols at each of the n orientations of the layout produces n representations of the orientation codeword, each representation comprising a different cyclic shift of the orientation codeword and being indicative of the degree of rotation of the layout, and wherein the orientation codeword is fault tolerant.

44. The system of any one of claims 24, 27 or 28 wherein the coded data is disposed on the interface surface in accordance with a layout, the layout having at least order n rotational symmetry, where n is at least two, the layout including n identical sub-layouts rotated $1/n$ revolutions apart about a center of rotational symmetry of the layout, the coded data disposed in accordance with each sub-layout including rotation-indicating data that distinguishes the rotation of that sub-layout from the rotation of at least one other sub-layout within the layout.

45. The system of any one of claims 24, 27 or 28 wherein the coded data is disposed on the interface surface in accordance with a layout having n -fold rotational symmetry, where n is at least two, the layout including n identical first sub-layouts rotated $1/n$ revolutions apart about a center of rotational symmetry of the layout, the coded data disposed in accordance with each first sub-layout including rotation-indicating data that distinguishes the rotation of that first sub-layout from the rotation of at least one other first sub-layout within the layout.

46. An interactive product item adapted for interaction with a user via a sensing device and a computer system, and comprising:

a product item having an identity;

an interface surface associated with the product item and having disposed thereon information relating to the product item and coded data indicative of the identity of the product item.

47. The interactive product item of claim 46 wherein the coded data is further indicative of a plurality of reference points of the interface surface.

48. The interactive product item of claim 46 or claim 47 wherein the coded data is further indicative of a parameter of an interaction.

49. The interactive product item of claim 46 wherein the interface surface is selected from the group comprising:

(a) a label;

(b) a package; and

(c) a surface of the product item itself.

50. The interactive product item of claim 49 wherein the coded data is disposed on the interface surface in accordance with a layout, the layout having at least order n rotational symmetry, where n is at least two, the layout encoding an orientation codeword comprising a sequence of an integer multiple m of n symbols, where m is one or more, each encoded symbol being distributed at n locations about a center of rotational symmetry of the layout such that decoding the symbols at each of the n orientations of the layout produces n representations of the orientation codeword, each representation comprising a different cyclic shift of the orientation codeword and being indicative of the degree of rotation of the layout, and wherein the orientation codeword is fault tolerant.

51. The interactive product item of claim 49 wherein the coded data is disposed on the interface surface in accordance with a layout, the layout having at least order n rotational symmetry, where n is at least two, the layout including n identical sub-layouts rotated $1/n$ revolutions apart about a center of rotational symmetry of the layout, the coded data disposed in accordance with each sub-layout including rotation-indicating data that distinguishes the rotation of that sub-layout from the rotation of at least one other sub-layout within the layout.

52. The interactive product item of claim 49 wherein the coded data is disposed on the interface surface in accordance with a layout having n-fold rotational symmetry, where n is at least two, the layout including n identical first sub-layouts rotated $1/n$ revolutions apart about a center of rotational symmetry of the layout, the coded data disposed in accordance with each first sub-layout including rotation-indicating data that distinguishes the rotation of that first sub-layout from the rotation of at least one other first sub-layout within the layout.

ABSTRACT

A method of facilitating an interaction between a user and a product item, the product item having an identity and the method including a number of steps. The first step involves providing the user with an interface surface associated with the product item and containing information relating to the product item. The interface surface includes coded data indicative of the identity of the product item and of a number of reference points of the interface surface. The next step involves receiving, in a computer system, indicating data from a sensing device regarding the identity of the product item and a position of the sensing device relative to the interface surface. The sensing device, when placed in an operative position relative to the interface surface, sensing the coded data in the vicinity of the sensing device and generating the indicating data using the sensed coded data. The next step involves facilitating, in the computer system and with reference to the indicating data, the interaction between the user and the product item.




FIG. 1




FIG. 2




FIG. 3




FIG. 4

5/47

17 ~

706 ~

726 ~

0G	4B	2K	4C	0C	1A	4G	1B	4K	1C	4O
1H	2M	4D	2L	1D	2H	3Z	2H	3M	2D	3L
2N	4E	2G	4F	2N	3Z	4F	3M	4F	2D	2A
3C	4D	2F	4E	3C	4D	2I	3E	4I	1F	3L
4A	2I	2N	4B	3G	3B	2J	3E	2I	4N	2A
5	3N	4M	2G	3C	3B	4H	4I	2E	3H	4M
6	4C	2L	3G	4D	4L	2K	2G	1N	4M	2C
7	2B	4K	2C	1O	3A	4I	2F	2G	1J	1K
8	1K	2M	2C	0	3A	4L	3H	3M	2B	2B
9	2G	2B	2C	1K	4A	2I	3E	2I	4N	2G
10	2B	2C	2C	1O	1O	2J	3E	2I	4N	2B
11	2C	2C	2C	1K	1K	2J	3E	2I	4N	2C
12	2K	2M	2C	1O	1O	2J	3E	2I	4N	2K
13	2L	2G	2C	1K	1K	2J	3E	2I	4N	2L
14	2H	2M	2C	1O	1O	2J	3E	2I	4N	2H
15	2G	2B	2C	1K	1K	2J	3E	2I	4N	2G
16	2F	2C	2C	1O	1O	2J	3E	2I	4N	2F
17	2E	2C	2C	1K	1K	2J	3E	2I	4N	2E
18	2D	2C	2C	1O	1O	2J	3E	2I	4N	2D
19	2C	2C	2C	1K	1K	2J	3E	2I	4N	2C
20	2B	2C	2C	1O	1O	2J	3E	2I	4N	2B
21	2A	2C	2C	1K	1K	2J	3E	2I	4N	2A
22	2G	2B	2C	1O	1O	2J	3E	2I	4N	2G
23	2F	2C	2C	1K	1K	2J	3E	2I	4N	2F
24	2E	2C	2C	1O	1O	2J	3E	2I	4N	2E
25	2D	2C	2C	1K	1K	2J	3E	2I	4N	2D
26	2C	2C	2C	1O	1O	2J	3E	2I	4N	2C
27	2B	2C	2C	1K	1K	2J	3E	2I	4N	2B
28	2A	2C	2C	1O	1O	2J	3E	2I	4N	2A
29	2G	2B	2C	1K	1K	2J	3E	2I	4N	2G
30	2F	2C	2C	1O	1O	2J	3E	2I	4N	2F
31	2E	2C	2C	1K	1K	2J	3E	2I	4N	2E
32	2D	2C	2C	1O	1O	2J	3E	2I	4N	2D
33	2C	2C	2C	1K	1K	2J	3E	2I	4N	2C
34	2B	2C	2C	1O	1O	2J	3E	2I	4N	2B
35	2A	2C	2C	1K	1K	2J	3E	2I	4N	2A
36	2G	2B	2C	1K	1K	2J	3E	2I	4N	2G
37	2F	2C	2C	1O	1O	2J	3E	2I	4N	2F
38	2E	2C	2C	1K	1K	2J	3E	2I	4N	2E
39	2D	2C	2C	1O	1O	2J	3E	2I	4N	2D
40	2C	2C	2C	1K	1K	2J	3E	2I	4N	2C
41	2B	2C	2C	1O	1O	2J	3E	2I	4N	2B
42	2A	2C	2C	1K	1K	2J	3E	2I	4N	2A
43	2G	2B	2C	1K	1K	2J	3E	2I	4N	2G
44	2F	2C	2C	1O	1O	2J	3E	2I	4N	2F
45	2E	2C	2C	1K	1K	2J	3E	2I	4N	2E
46	2D	2C	2C	1O	1O	2J	3E	2I	4N	2D
47	2C	2C	2C	1K	1K	2J	3E	2I	4N	2C
48	2B	2C	2C	1O	1O	2J	3E	2I	4N	2B
49	2A	2C	2C	1K	1K	2J	3E	2I	4N	2A
50	2G	2B	2C	1K	1K	2J	3E	2I	4N	2G
51	2F	2C	2C	1O	1O	2J	3E	2I	4N	2F
52	2E	2C	2C	1K	1K	2J	3E	2I	4N	2E
53	2D	2C	2C	1O	1O	2J	3E	2I	4N	2D
54	2C	2C	2C	1K	1K	2J	3E	2I	4N	2C
55	2B	2C	2C	1O	1O	2J	3E	2I	4N	2B
56	2A	2C	2C	1K	1K	2J	3E	2I	4N	2A
57	2G	2B	2C	1K	1K	2J	3E	2I	4N	2G
58	2F	2C	2C	1O	1O	2J	3E	2I	4N	2F
59	2E	2C	2C	1K	1K	2J	3E	2I	4N	2E
60	2D	2C	2C	1O	1O	2J	3E	2I	4N	2D
61	2C	2C	2C	1K	1K	2J	3E	2I	4N	2C
62	2B	2C	2C	1O	1O	2J	3E	2I	4N	2B
63	2A	2C	2C	1K	1K	2J	3E	2I	4N	2A
64	2G	2B	2C	1K	1K	2J	3E	2I	4N	2G
65	2F	2C	2C	1O	1O	2J	3E	2I	4N	2F
66	2E	2C	2C	1K	1K	2J	3E	2I	4N	2E
67	2D	2C	2C	1O	1O	2J	3E	2I	4N	2D
68	2C	2C	2C	1K	1K	2J	3E	2I	4N	2C
69	2B	2C	2C	1O	1O	2J	3E	2I	4N	2B
70	2A	2C	2C	1K	1K	2J	3E	2I	4N	2A
71	2G	2B	2C	1K	1K	2J	3E	2I	4N	2G
72	2F	2C	2C	1O	1O	2J	3E	2I	4N	2F
73	2E	2C	2C	1K	1K	2J	3E	2I	4N	2E
74	2D	2C	2C	1O	1O	2J	3E	2I	4N	2D
75	2C	2C	2C	1K	1K	2J	3E	2I	4N	2C
76	2B	2C	2C	1O	1O	2J	3E	2I	4N	2B
77	2A	2C	2C	1K	1K	2J	3E	2I	4N	2A
78	2G	2B	2C	1K	1K	2J	3E	2I	4N	2G
79	2F	2C	2C	1O	1O	2J	3E	2I	4N	2F
80	2E	2C	2C	1K	1K	2J	3E	2I	4N	2E
81	2D	2C	2C	1O	1O	2J	3E	2I	4N	2D
82	2C	2C	2C	1K	1K	2J	3E	2I	4N	2C
83	2B	2C	2C	1O	1O	2J	3E	2I	4N	2B
84	2A	2C	2C	1K	1K	2J	3E	2I	4N	2A
85	2G	2B	2C	1K	1K	2J	3E	2I	4N	2G
86	2F	2C	2C	1O	1O	2J	3E	2I	4N	2F
87	2E	2C	2C	1K	1K	2J	3E	2I	4N	2E
88	2D	2C	2C	1O	1O	2J	3E	2I	4N	2D
89	2C	2C	2C	1K	1K	2J	3E	2I	4N	2C
90	2B	2C	2C	1O	1O	2J	3E	2I	4N	2B
91	2A	2C	2C	1K	1K	2J	3E	2I	4N	2A
92	2G	2B	2C	1K	1K	2J	3E	2I	4N	2G
93	2F	2C	2C	1O	1O	2J	3E	2I	4N	2F
94	2E	2C	2C	1K	1K	2J	3E	2I	4N	2E
95	2D	2C	2C	1O	1O	2J	3E	2I	4N	2D
96	2C	2C	2C	1K	1K	2J	3E	2I	4N	2C
97	2B	2C	2C	1O	1O	2J	3E	2I	4N	2B
98	2A	2C	2C	1K	1K	2J	3E	2I	4N	2A
99	2G	2B	2C	1K	1K	2J	3E	2I	4N	2G
100	2F	2C	2C	1O	1O	2J	3E	2I	4N	2F
101	2E	2C	2C	1K	1K	2J	3E	2I	4N	2E
102	2D	2C	2C	1O	1O	2J	3E	2I	4N	2D
103	2C	2C	2C	1K	1K	2J	3E	2I	4N	2C
104	2B	2C	2C	1O	1O	2J	3E	2I	4N	2B
105	2A	2C	2C	1K	1K	2J	3E	2I	4N	2A
106	2G	2B	2C	1K	1K	2J	3E	2I	4N	2G
107	2F	2C	2C	1O	1O	2J	3E	2I	4N	2F
108	2E	2C	2C	1K	1K	2J	3E	2I	4N	2E
109	2D	2C	2C	1O	1O	2J	3E	2I	4N	2D
110	2C	2C	2C	1K	1K	2J	3E	2I	4N	2C
111	2B	2C	2C	1O	1O	2J	3E	2I	4N	2B
112	2A	2C	2C	1K	1K	2J	3E	2I	4N	2A
113	2G	2B	2C	1K	1K	2J	3E	2I	4N	2G
114	2F	2C	2C	1O	1O	2J	3E	2I	4N	2F
115	2E	2C	2C	1K	1K	2J	3E	2I	4N	2E
116	2D	2C	2C	1O	1O	2J	3E	2I	4N	2D
117	2C	2C	2C	1K	1K	2J	3E	2I	4N	2C
118	2B	2C	2C	1O	1O	2J	3E	2I	4N	2B
119	2A	2C	2C	1K	1K	2J	3E	2I	4N	2A
120	2G	2B	2C	1K	1K	2J	3E	2I	4N	2G
121	2F	2C	2C	1O	1O	2J	3E	2I	4N	2F
122	2E	2C	2C	1K	1K	2J	3E	2I	4N	2E
123	2D	2C	2C	1O	1O	2J	3E	2I	4N	2D
124	2C	2C	2C	1K	1K	2J	3E	2I	4N	2C
125	2B	2C	2C	1O	1O	2J	3E	2I	4N	2B
126	2A	2C	2C	1K	1K	2J	3E	2I	4N	2A
127	2G	2B	2C	1K	1K	2J	3E	2I	4N	2G
128	2F	2C	2C	1O	1O	2J	3E	2I	4N	2F
129	2E	2C	2C	1K	1K	2J	3E	2I	4N	2E
130	2D	2C	2C	1O	1O	2J	3E	2I	4N	2D
131	2C	2C	2C	1K	1K	2J	3E	2I	4N	2C
132	2B	2C	2C	1O	1O	2J	3E	2I	4N	2B
133	2A	2C	2C	1K	1K	2J	3E	2I	4N	2A
134	2G	2B	2C	1K	1K	2J	3E	2I	4N	2G
135	2F									

6/47




FIG. 5c

7/47




FIG. 6




FIG. 7

9/47




FIG. 8




FIG. 9




FIG. 10

12/47




FIG. 11




FIG. 12




FIG. 12a




FIG. 13




FIG. 14




FIG. 15




FIG. 16




FIG. 17

20/47




FIG. 18




FIG. 19




FIG. 20




FIG. 21

23/47




FIG. 22




FIG. 23




FIG. 24




FIG. 25

26/47




FIG. 26




FIG. 27




FIG. 28

28/47




FIG. 29




FIG. 30




FIG. 31




FIG. 32

30/47




FIG. 33




FIG. 34




FIG. 35




FIG. 36




FIG. 37




FIG. 38




FIG. 38a




FIG. 39




FIG. 40




FIG. 41




FIG. 42




FIG. 43




FIG. 44




FIG. 45




FIG. 46




FIG. 47




FIG. 48

41/47




FIG. 49

42/47




FIG. 50

43/47




FIG. 51




FIG. 52




FIG. 53





FIG. 54



47/47




FIG. 56




FIG. 57

**This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

BLACK BORDERS

IMAGE CUT OFF AT TOP, BOTTOM OR SIDES

FADED TEXT OR DRAWING

BLURRED OR ILLEGIBLE TEXT OR DRAWING

SKEWED/SLANTED IMAGES

COLOR OR BLACK AND WHITE PHOTOGRAPHS

GRAY SCALE DOCUMENTS

LINES OR MARKS ON ORIGINAL DOCUMENT

REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY

OTHER: _____

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.